Volume VI: Philosophy of Physics

Ahmed Ali Saleh[ADAM]

October 5, 2025

Contents

I Orientation and Thesis **1**

General Introduction to the Volume **3**

1 Why Abstraction Misleads (and How to Tame It) **5**

1.1 Thesis . . . 5

1.1.1 Mental vs Physical . . . 6

1.1.2 What Counts as Real . . . 7

1.2 Four Control Principles . . . 8

1.2.1 Quick Exemplars . . . 9

1.3 Roadmap . . . 10

1.3.1 Acceptance Metrics . . . 11

2 Early Lexicon and Units **15**

2.1 Symbols and Meanings . . . 15

2.1.1 Physical Proxies . . . 16

2.2 Observables and Units . . . 18

2.2.1 Finite, Unit-Bearing . . . 18

2.3 Gauge-Invariant Descriptors . . . 20

2.3.1 Reality via Invariance . . . 20

3 Method and Style of Claims **25**

3.1 Domain of Validity . . . 25

3.1.1 Declaring DoV . . . 25

3.1.2 End-of-Validity . . . 26

3.2 Idealization Audits . . . 28

3.2.1 Stability at Finite . . . 28

3.3 Math-to-Measurement Pipeline . . . 31

3.3.1 Model → Pointers . . . 32

4 The Spectrum-First Principle **35**

4.1 Statement and Rationale . . . 35

4.2 From Spectral Fields to Particles (Solitons) 36
4.3 Spectral Immersion and the Classical Limit 36
4.4 Device Proxies and Falsifiable Signatures 37
4.5 Relation to Gates and DoV 37

II The Abstract and the Real 39

5 Operational Semantics of Physical Subject 43
5.1 Definitions 43
5.1.1 Physical vs Mathematical 43
5.2 Gauge Realism 46
5.2.1 Invariance Tests 46
5.3 Bridging Maps 49
5.3.1 Delayed Measurability 50

III Infinities, Singularities, End-of-Validity 53

6 Infinity as a Limit, Not a Value 57
6.1 Legitimate Uses of ∞ 57
6.1.1 Asymptotics, Packets 58
6.2 Stability-at-Finite Tests 60
6.2.1 Protocols 61
6.3 Case Studies 63
6.3.1 UV, Plane Waves 63

7 Anti-Singularity 67
7.1 Coordinate vs Curvature 67
7.1.1 Invariants 69
7.2 End-of-Validity Flags 70
7.2.1 Effective Replacements 71
7.3 Decision Protocol 73
7.3.1 Three Steps 74

IV Measurement, Open Dynamics, and Probability 77

8 From Model to Pointer 81
8.1 Open Systems and GKSL 81
8.1.1 Causality and CP 82
8.2 Pointer Observables 85

8.2.1 Calibration, Frequencies . . . 85
8.3 Noise and Timescales . . . 87
8.3.1 Relaxation vs Decoherence . . . 87

9 Probability Without Postulates 91
9.1 Born as Stable Energy Shares . . . 91
9.1.1 Long-Time Behavior . . . 91
9.2 Competing Accounts . . . 94
9.2.1 Assumptions Exposed . . . 94
9.3 Empirical Discriminants . . . 96
9.3.1 What Could Differ . . . 96

V Laws as Stable Regularities 101

10 Stationarity, Symmetries, Coarse-Graining 105
10.1 Patterns that Survive . . . 105
10.1.1 RG Intuition . . . 106
10.2 Background Independence . . . 108
10.2.1 What It Forbids/Allows . . . 108
10.3 Measuring Law Strength . . . 111
10.3.1 Across Scales . . . 111

VI Beginnings and the Cosmic Narrative 115

11 The Principle of Physical Beginning 119
11.1 Statement and Rationale . . . 119
11.1.1 No “Nothing” . . . 119
11.2 First Describable System . . . 121
11.2.1 Seeds of g_{eff} . . . 121
11.2.2 Prototype constructions . . . 123

12 Acceptance Checklist for a Beginning 125
12.1 Six Conditions . . . 125
12.1.1 Failure Modes . . . 127
12.2 Scope of Physics . . . 128
12.2.1 Inside/Outside . . . 129
12.3 Rules for Abstraction . . . 130
12.3.1 Pointer Supremacy . . . 131

13 Early-to-Late Bridges **133**
13.1 Bridging Scenarios . . . 133
13.1.1 Generic Signatures . . . 134
13.2 Toy Pipelines . . . 135
13.2.1 Pipeline A: Diffusive seed → spectral knee . . . 136
13.2.2 Pipeline B: Directional seed → birefringence slope . . . 136
13.2.3 Pipeline C: Open quantum seed → energy–share vector . . . 137
13.2.4 Late-Time Pointers . . . 137

VII Consistency Gates as Epistemic Virtues **139**

14 Defining and Testing S/G/TQ/D **143**
14.1 Spectral Safety . . . 143
14.1.1 Operator Tests . . . 144
14.2 Causality within g_{eff} . . . 146
14.2.1 Cones, No-Signalling . . . 146
14.3 Complete Positivity . . . 149
14.3.1 Subsystem Consistency . . . 150
14.4 Non-Negative Dissipation . . . 152
14.4.1 Entropy Production . . . 153

15 Gate-Based Audits **155**
15.1 Three Mini Case Studies . . . 155
15.1.1 GR/QFT/Media . . . 155
15.2 Writing the Gate Report . . . 157
15.2.1 Templates . . . 158

VIII Virtues and Limits of Mathematics **161**

16 Proofs, Domains, Ethics of Modelling **165**
16.1 What Proofs Can/Can't Do . . . 165
16.1.1 Beyond Measurability . . . 167
16.2 Idealization without Idolatry . . . 168
16.2.1 Reporting Assumptions . . . 169
16.3 Style Guide for Equations . . . 170
16.3.1 The Five Rules . . . 170

IX Objections and Replies 173

17 Instrumentalism, Structural Realism, Born 177

17.1 Just Instrumentalism? . . . 177
17.1.1 Gauge Realism . . . 178
17.2 Only Structure Is Real? . . . 179
17.2.1 Anchored Patterns . . . 180
17.3 Is Born Assumed? . . . 182
17.3.1 Pipeline Derivation . . . 183

18 What About Gauge? 187

18.1 Reality vs Description . . . 187
18.1.1 Anti-Reification . . . 189
18.2 Pragmatics . . . 190
18.2.1 Writing Non-Reifying Claims . . . 190

X Program, Notes, Appendices 193

19 A Research Program 197

19.1 From Singularity to DoV . . . 197
19.1.1 Effective Replacements . . . 198
19.2 Gauge-Invariant Existence . . . 199
19.2.1 QFT/GR . . . 200
19.3 Probability as Dissipation . . . 202
19.3.1 Open Models . . . 203
19.4 Internal Time . . . 205
19.4.1 Difference Observables . . . 206

20 Notation and Conventions 209

20.1 Typography and naming . . . 209
20.2 Indices, tensors, and sign conventions . . . 210
20.3 Units, dimensions, constants . . . 210
20.4 Fourier/Laplace and spectral conventions . . . 211
20.5 Open dynamics and GKSL sign conventions . . . 211
20.6 Probability and statistics (reporting) . . . 211
20.7 Domains of validity (DoV), gates, and badges . . . 212
20.8 Spectrum-first descriptors (object language policy) . . . 212
20.9 Figures, tables, and captions . . . 212
20.10 Citations and reproducibility (practical) . . . 212
20.11 Worked micro-templates . . . 213

21 Mini-Derivations and Worked Examples **215**
21.1 Mini-derivation 1 — PSD non-negativity at finite resolution 215
21.2 Mini-derivation 2 — Born as long-time dissipative shares 215
21.3 Mini-derivation 3 — Stability-at-finite acceptance test 216
21.4 Worked example A — UV observable with a hard cutoff 216
21.5 Worked example B — Front vs group velocity (no–signalling) 217
21.6 Worked example C — Resonance pole extraction (QFT) 217
21.7 Worked example D — GR ringdown as invariant spectroscopy 218
21.8 Worked example E — Internal time via frequency ratios 218
21.9 Worked example F — Spectrum-first soliton audit (NLSE proxy) 218
21.10 Worked example G — CP of GKSL maps (Spohn inequality sketch) . . . 219
21.11 Worked example H — Coordinate artefact audit (GR horizon) 219

22 Gate Checklists and Templates **221**
22.1 One–page overview (what each gate asks for) 221
22.2 Operational tests by gate 221
22.3 Gate cards and report templates 223
22.4 Filled examples (what a passing card looks like) 223
22.5 Author templates by domain (paste and fill) 224

Part I

Orientation and Thesis

General Introduction to the Volume

This volume develops a disciplined *philosophy of physics* fit for practice. Its central thesis is simple to state and hard to honor: *mental absolutes are not physical quantities.* Infinity, perfect continua, exact symmetries, and other ideal forms are extraordinary tools of thought; they are not themselves objects in nature. A physical claim earns its status only when it can be connected, in principle, to measurement via a causal pipeline that survives coarse-graining and respects basic consistency gates: spectral safety, causality, complete positivity, and non-negative dissipation.

Three methodological moves structure the book:

1. **From abstraction to operation.** We attach *operational semantics* to all key notions. A quantity is "physical" only if there exists a conceived measurement procedure yielding a *finite, unit-bearing* outcome that remains stable under approximations below device resolution.

2. **From laws to stability.** Laws are treated as *stable regularities*—patterns that persist under symmetry and coarse-graining—rather than as a pre-given stage. This reframing clarifies what "background independence" forbids and allows.

3. **From postulates to pipelines.** Wherever possible, postulates are replaced by derivation pipelines. In particular, probabilities are read as *stable energy shares* emerging in open dynamics (e.g. via Gorini–Kossakowski–Sudarshan–Lindblad (GKSL) generators), checked against long-time pointer behavior.

4. **From objects to spectra.** We adopt a *spectrum-first* stance: modelling begins with spectral fields and mode shares; "particles/objects" are stabilized spectral organizations (packets, solitons, poles) admitted only on a declared DoV and with a device-level bridge.

title=How to read spectral claims

Primitives: power spectral densities $S_{YY}(\omega)$, calibrated mode shares $\{p_m\}$ (non-negative, sum to 1), peak sets/Q-factors, and scattering poles/residues (with units). **Stability at finite:** report window/band/time dependence and a tolerance; show that the result is stable under coarse-graining.

Gate line: S (non-negative spectra), G (front/cone consistent with g_{eff}), TQ (CP under reductions), D (non-negative entropy/linewidth).
Object language: only after spectral anchoring—state how the spectral data license it.

Criterion 0.1: Acceptance Metric

A claim is acceptable to physics when (i) it is framed in gauge-invariant, finite descriptors; (ii) it passes the S/G/TQ/D gates (spectral safety, causality within g_{eff}, Complete Positivity (CP), non-negative dissipation); and (iii) it offers a clear bridge from model descriptors to pointer observables. Otherwise, the claim is *abstract* or *metaphysical* rather than physical.

The chapters that follow operationalize these ideas. Part I clarifies the central thesis and sets the style of claims. Later parts address infinities and singularities (as signs of end-of-validity), measurement and probability (without extra postulates), beginnings and the cosmic narrative (without invoking "nothing" as a physical subject), and the gate-based audits that keep us honest.

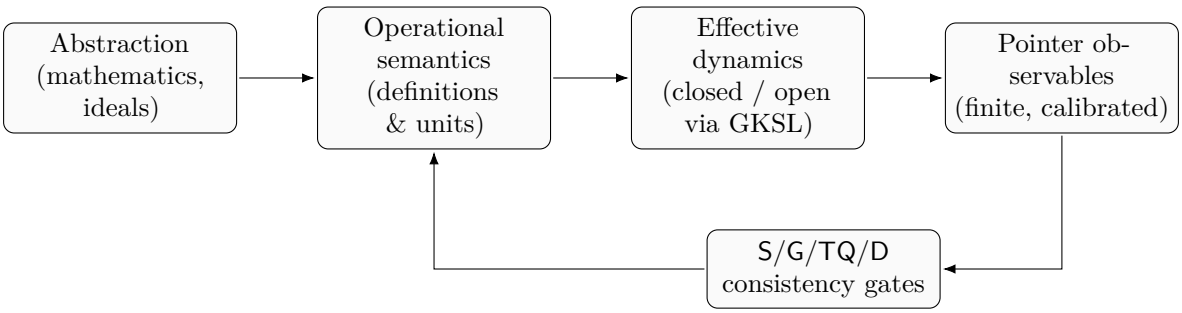

Figure 0.1: Math-to-measurement pipeline with feedback through the S/G/TQ/D gates.

As a discipline for claims about *objects*, we adopt a spectrum-first policy: invariants in the spectral domain anchor the claim, a stability-at-finite audit protects the limit, and only then do the S/G/TQ/D gates license the object language, as summarized in Figure 0.2.

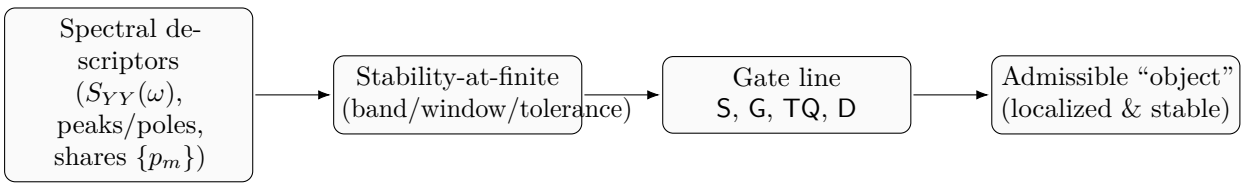

Figure 0.2: "Spectrum first": object language is licensed by a spectral anchor plus stability-at-finite and gate checks.

Chapter 1

Why Abstraction Misleads (and How to Tame It)

1.1 Thesis

Principle 1.1: No Reification of Abstraction

Do not promote abstract devices (e.g. limits, ideal symmetries, coordinate choices) to physical entities. In physics, existence is tied to *measurable, finite, unit-bearing* outcomes that remain stable under coarse-graining and that can be situated within causal dynamics.

Abstraction is indispensable; it compresses, generalizes, and reveals structure. But abstraction also tempts us to treat calculational conveniences as if they were things in the world. Plane waves, point particles, exact continua, and ∞ are powerful *limits*; they are not values to be read off a device. This chapter lays down the discipline that prevents such slippage, by (i) distinguishing mental from physical subjects, and (ii) specifying what, in this book, *counts as real*.

Working Attitude

We retain all the power of mathematics while refusing to idolize it. Every equation arrives with a declared domain of validity, units, and a pointer-bridge. Every idealization is audited for stability at finite approximation. Every gauge choice is flagged as descriptive, not ontic.

1.1.1 Mental vs Physical

Definition 1.1: Physical Subject

A subject is *physical* if there exists a conceived measurement procedure that, in principle, yields a *finite* outcome with *units*, is *causally* admissible within an effective cone g_{eff}, and remains *stable* under coarse-graining below device resolution.

Definition 1.2: Mathematical Subject

A subject is *mathematical* if it is defined purely within a formal structure (e.g. as a limit, idealization, or coordinate artifact) with no direct measurement procedure attached. Such entities guide reasoning and approximations but are not, by themselves, physical.

The distinction is *operational*, not metaphysical. It does not deny the reality of patterns discoverable by mathematics; it demands a *measurement-facing* interpretation. Consider:

- *Infinity (∞).* Accepted *as a limit* when predictions are *stable at finite* truncations; rejected as a *value*.
- *Plane waves.* Excellent asymptotic templates; physical radiation arrives as finite-duration packets with finite energy flux and bandwidth.
- *Point particles.* Useful effective descriptions within a domain of validity; outside it, extended structure or new degrees of freedom must replace the fiction.

Principle 1.2: Infinity-as-Limit.

Infinity is admissible only as a limiting device; a result is physically acceptable when it is demonstrably *stable* at finite approximation.

A practical stability check is to demand, for a sequence of finite approximants (Q_N),

$$\exists\, N_0,\ \forall N \geq N_0 : \quad \frac{|Q_{N+1} - Q_N|}{\max\{1, |Q_N|\}} < \varepsilon, \tag{1.1}$$

for a tolerance ε appropriate to the device and context. Passing (1.1) licenses the use of the limiting idealization in *this* context; failing it signals end-of-validity or the need for a different effective description.

Principle 1.3: Anti-Singularity

A divergence in a non-observable or a blow-up traceable to coordinates announces the *end of validity* of a model, not a new thing in nature. The cure is reparametrization

or an improved effective theory, not reification of the singular behavior.

1.1.2 What Counts as Real

Criterion 1.1: Gauge Reality

What remains invariant under admissible gauge/coordinate transformations is *real*. Quantities that depend on descriptive choices are not candidates for ontology; they are bookkeeping devices.

This criterion is not a mere slogan; it is a workflow requirement:

1. **State in invariants.** Formulate claims in terms of gauge- or coordinate-invariant descriptors (e.g. eigenvalues, curvature scalars, spectral data).
2. **Give the bridge.** Provide a *pointer-bridge*: a map from descriptors to device readings (calibration curves, long-time frequencies, energy-share readouts).
3. **Pass the gates.** Check S/G/TQ/D: no negative spectra for observables; causality within g_{eff}; CP for subsystems; non-negative dissipation defining an arrow of time.

title=From Postulates to Pipelines

Instead of positing probabilities, we *derive* stable frequencies from open-system dynamics (e.g. GKSL generators) and long-time pointer behavior. Instead of positing backgrounds, we *exhibit* effective cones g_{eff} emerging with the descriptors that survive coarse-graining.

Two frequent temptations follow from ignoring the criterion:

Reifying coordinates. Treating a convenient chart as if it were the world can turn removable artifacts into dramatic "effects". The cure is to rewrite in invariants and test against devices.

Idolizing ideals. Taking an ideal limit to literally describe nature yields "infinite energies", "singular densities", or "exact symmetries" where none can be measured. The cure is stability-at-finite and explicit domains of validity.

title=Style Guide for Claims (Chapter Tag)

Every claim in this chapter (and the volume) should come with: (i) units and domain of validity; (ii) an indication of gauge status (invariant vs choice); (iii) the intended pointer-bridge; and (iv) a brief S/G/TQ/D checklist. This is not bureaucracy; it is how we keep abstraction honest.

Gate 1.1: S/G/TQ/D Audit for Chapter 1

S (Spectral Safety): Claims are stated without negative spectra for observables.
G (Causality within g_{eff}): All measurement procedures are assumed to lie within admissible effective cones.
TQ (Complete Positivity): Open-system evolutions are treated with CP generators (e.g. GKSL).
D (Non-Negative Dissipation): Irreversibility appears only with non-negative dissipation; no arrow without activity.

1.2 Four Control Principles

The volume is governed by four control principles. They are not slogans; they are *filters* that keep abstraction honest, force claims to carry units and a domain of validity, and anchor all talk to pointer observables and causal dynamics. Each principle below includes its operational reading.

Principle 1.4: No Background Stage (as a discipline).

Do not assume a pre-given arena (global frame, absolute time, fixed metric) independent of the physical descriptors. Geometry and causal structure are *effective* and must be read from the same descriptors that enter the dynamics (e.g. an emergent effective cone g_{eff}).

Rule 1.1:

Formulate claims in gauge/coordinate invariants. If a statement changes under a reparametrization or gauge choice, it is descriptive, not ontic. When cones or metrics are invoked, exhibit how they emerge from the medium/descriptors used in the model.

Principle 1.5: Infinity-as-Limit.

∞ is admissible only as a limiting device. A result is physically acceptable when it is demonstrably *stable at finite* approximation (below device resolution); ∞ is never a measured value.

Rule 1.2:

Provide a stability-at-finite audit: for finite approximants (Q_N), verify convergence in the sense of

$$\exists N_0 : \ \forall N \geq N_0, \quad \frac{|Q_{N+1} - Q_N|}{\max\{1, |Q_N|\}} < \varepsilon$$

for a tolerance ε justified by calibration. Failure flags end-of-validity or the need for a different effective description.

Principle 1.6: Anti-Singularity.

A divergence in non-invariant or non-observable quantities, or a blow-up traceable to coordinates, announces *end of validity* for the model—not a new physical entity. The cure is reparametrization or an improved effective theory.

Rule 1.3:

(i) Check whether the divergence disappears in invariants; (ii) test stability-at-finite for observables; (iii) if both fail, declare a domain boundary and replace the model (e.g. by coarse-grained or UV-completed dynamics).

Principle 1.7: Derive, Don't Postulate.

Replace free-standing postulates with derivation pipelines from model descriptors to pointer observables via effective (often open) dynamics. Read probabilities as *stable energy shares* in long-time pointer behavior, e.g. under GKSL generators with non-negative rates.

Rule 1.4:

Exhibit a map

$$\text{model descriptors} \longrightarrow \text{effective dynamics (closed/open)} \longrightarrow \text{pointer observables},$$

and pass the S/G/TQ/D gates: spectral safety, causality within g_{eff}, complete positivity, and non-negative dissipation.

1.2.1 Quick Exemplars

No Background Stage.

- *GR (discipline).* Schwarzschild coordinates create a spurious "singularity" at $r = 2M$; curvature invariants remain finite there. Geometry is read from invariants, not from a chart. Cosmological metrics are solutions, not a stage assumed in advance.

- *Effective media.* In anisotropic media, directional order (our descriptor analogous to χ^μ) selects an effective cone g_{eff}; signal speeds depend on state and scale, not on a fixed background.

Infinity-as-Limit.

- *Analysis.* Plane waves and Dirac deltas are ideal limits. Physical radiation arrives as finite-energy packets; acceptance requires a stability-at-finite check against device bandwidth.
- *QFT.* UV divergences signal sensitivity to unresolved scales. Renormalization replaces them by running, scale-dependent couplings; acceptance rests on regulator independence within quoted uncertainty.

Anti-Singularity.

- *GR.* The blow-up at $r = 0$ in Schwarzschild reflects true curvature growth and marks end-of-validity (missing microphysics). At $r = 2M$, the divergence is coordinate and removable.
- *Field theory.* Landau poles or non-physical negative norms are refusal signals for the model (ghosts/instabilities), not new particles.

Derive, Don't Postulate.

- *Measurement pipeline.* Under GKSL evolution,

$$\dot{\rho} = -i[H, \rho] + \sum_k \gamma_k \Big(L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\} \Big), \qquad \gamma_k \geq 0,$$

  long-time pointer frequencies stabilize and can be interpreted as energy-share probabilities without an extra postulate, provided S/G/TQ/D hold.
- *Constants as effective.* "Couplings" and characteristic speeds are read as stable patterns under coarse-graining; they are outputs of the same descriptors that define g_{eff}.

Checklist

One-page audit. For any chapter claim: (1) invariant form; (2) units & domain of validity; (3) stability-at-finite if a limit is used; (4) pointer bridge; (5) S/G/TQ/D checklist. Failing any item downgrades the claim to abstract/metaphysical until repaired.

1.3 Roadmap

This volume is designed as a *working* philosophy of physics. It moves from conceptual hygiene to operational practice and back again, with each part delivering something the rest can use.

title=Reading Map (What Each Part Gives You)

P1: Orientation & Thesis. The core stance: mental absolutes are not physical quantities; claims must carry units, a domain of validity (DoV), and a pointer-bridge, and must pass S/G/TQ/D gates.

P2: The Abstract and the Real. Operational semantics for *physical subjects*; gauge realism as the reality criterion; how to write non-reifying claims.

P3: Infinities, Singularities, End-of-Validity. How to use ∞ as a limit; how to recognize coordinate vs. curvature singularities; a three-step decision protocol when blow-ups appear.

P4: Measurement, Open Dynamics, Probability. Model $\to$ effective (often open) dynamics $\to$ pointer observables; deriving stable frequencies (Born) as long-time energy shares under GKSL.

P5: Laws as Stable Regularities. Laws as patterns that survive symmetry and coarse-graining; measuring "law strength" as stability across scales; background independence as discipline, not slogan.

P6: Beginnings & Cosmic Narrative. The *principle of physical beginning*: first describable systems in state space (not "nothing"); acceptance checklist and empirical bridge rules.

P7: Consistency Gates as Epistemic Virtues. Spectral safety, causality within g_{eff}, complete positivity, and non-negative dissipation as the standing audit.

P8: Virtues & Limits of Mathematics; Objections & Replies; Program. Style guide for equations and idealizations; replies to stock objections; short research notes you can publish.

The remainder of this section turns the map into a concrete *acceptance metric*—a checklist that every claim in the book (and in prospective follow-up papers) must satisfy.

1.3.1 Acceptance Metrics

Criterion 1.2: Spectrum-first addendum to acceptance

Any claim that uses *object/particle* language must supply a *spectral anchor*: (i) a calibrated descriptor (shares/PSD/peaks/poles) with units; (ii) a stability-at-finite audit over bands, windows, and acquisition time; (iii) an S/G/TQ/D gate line; and (iv) a pointer bridge (how the spectral quantity is read). Absent these, the claim is merely descriptive and should be rewritten in spectral terms.

Criterion 1.3: Acceptance Metric (Baseline).

A claim is *acceptable to physics* iff it is (i) expressed in gauge-invariant, *finite* descriptors with units; (ii) equipped with a declared domain of validity and calibration context; (iii) connected by a clear model → dynamics → pointer pipeline; and (iv) compliant with the four consistency gates S/G/TQ/D.

We now refine the baseline into actionable items. Each item has a pass/fail test; failing any Baseline item demotes the claim to abstract/metaphysical until repaired.

Checklist

Baseline Items (must-pass).

1. **Invariance & Units.** Statement is in gauge/coordinate invariants; every quantity carries physical units. *Test:* rewrite in invariants (e.g. eigen/spectral data, curvature scalars) and list units.
2. **DoV Declaration.** Domain of validity is stated with boundaries (scales, regimes, approximations). *Test:* explicit DoV sentence and boundary trigger (what would mark end-of-validity).
3. **Stability-at-Finite (when limits are used).** If an ideal limit or ∞ enters, provide a finite-approximation audit with tolerance ε justified by calibration:

$$\exists N_0 : \ \forall N \geq N_0, \quad \frac{|Q_{N+1} - Q_N|}{\max\{1, |Q_N|\}} < \varepsilon.$$

4. **Pipeline Shown.** Map from descriptors to effective dynamics (closed/open) to pointer observables, with the device timescales noted.
5. **S/G/TQ/D Compliance.**
   - **S** (Spectral safety): no negative spectra for observables.
   - **G** (Causality within g_{eff}): no superluminal signalling in the effective cone.
   - **TQ** (Complete positivity): subsystem consistency.
   - **D** (Non-negative dissipation): properly signed rates; entropy production non-negative on accessible timescales.

Rule 1.5:

- **Provisional** ✓ = Baseline passes.
- **Strong** ✓✓ = Baseline + quantitative error budget (noise, calibration, estima-

tor bias) and robustness under modest model variations (sensitivity analysis).

- **Publication-ready ✓✓✓** = Strong + falsifiable empirical signature (sign/scale/pattern) and a reproducible script or pseudocode for the stability/DoV audits.

Reviewer's Gate Report (one page). For each substantive claim or figure, attach a short gate report:

title=Gate Report Template

Claim: *(one sentence in invariants with units)*
DoV: *(regimes/scales; boundary trigger)*
Pipeline: *descriptors* $\to$ *dynamics* (H, L_k, *rates*) $\to$ *pointer*
S: *operator/spectrum check* **G:** *cone/no-signalling check*
TQ: *CP of reduced maps* **D:** *rate signs/entropy production*
Limits Audit: *stability-at-finite with* ε *& plot/table ref*
Outcome: *Provisional / Strong / Publication-ready*

Rejection Triggers (hard stops).

- Gauge/coordinate dependence disguised as ontology.
- Use of ∞ as a value, or lack of a stability-at-finite audit.
- Gate failure: ghosts/negative spectra; signalling outside g_{eff}; loss of complete positivity; negative dissipation on the reported timescale.
- No declared DoV or missing pointer-bridge.

Where This Section is Used. Every part of the volume inherits these metrics. In Section 1.2 they police abstraction; in the measurement chapters they certify the model $\to$ pointer pipeline; in the beginnings narrative they filter metaphysics from physics; and in the case studies they provide a uniform audit trail.

Chapter 2

Early Lexicon and Units

This chapter fixes a *working lexicon*. Our symbols are not metaphysical labels; they are *descriptors* tied to measurement ideas, units, and gauge-invariant content. Throughout the volume a statement earns physical status only when its descriptors are (i) finite and unit-bearing, (ii) gauge/coordinate invariant (or explicitly qualified as choices), and (iii) connected to device-level proxies.

Definition 2.1: Descriptor (working).

A descriptor is a model quantity intended to survive to the level of pointer observables. Concretely, it is specified with units, an anticipated calibration procedure, and a gauge-invariant reading (or a clear account of how a gauge choice is eliminated in the device output).

title=Design goal

Every symbol we introduce below comes with (a) a short physical meaning, (b) typical *proxies* for measurement, and (c) a note on invariance and domain-of-validity. This prevents "symbol drift" as the theory scales up.

2.1 Symbols and Meanings

We will repeatedly use five core descriptors for the quantum-aether medium and its effective behavior:

ρ **(rho) — Effective density / stiffness.**
Governs characteristic speeds and scale-setting in the medium (elastic/acoustic analogues). Conceptually, ρ modulates how "hard" it is for phase patterns to deform.

η **(eta) — Effective dissipation / viscosity.**
Encodes irreversible rates (relaxation, diffusion, linewidths). Non-negative η imple-

ments the arrow of time at the effective level.

ϕ **(phi) — Phase / order-parameter content.**
A phase-like matter descriptor capturing coherence, condensate weight, or order-parameter amplitude that transports on the background set by ρ.

T — **Internal clock (differences/rates only).**
Not a globally meaningful absolute time; only differences and local rates enter. T organizes temporal orderings recorded by the medium.

χ^μ **(chi–mu) — Directional order / transport bias.**
A vector-like descriptor that organizes directional transport and defines the effective causal cone g_{eff} (ASCII: `g_eff`).

Remark. These symbols are *not* fields to be reified in isolation; they are parts of a *descriptor tuple* whose gauge-invariant combinations feed the model → dynamics → pointer pipeline.

2.1.1 Physical Proxies

To keep the lexicon operational, each descriptor is paired with concrete measurement ideas. What follows are *typical* proxies; the later chapters specify device-level models and calibrations.

ρ — **density/stiffness.**

- **Proxies:** phase/group velocities of small-amplitude waves; slopes of stress–strain curves; elastic moduli analogues.
- **Units:** typically [energy · time^2/length^3] or device-specific elastic units.
- **Invariance:** read out from *ratios* of speeds or moduli that are gauge/coordinate invariant; report the domain of linear response.

η — **dissipation/viscosity.**

- **Proxies:** relaxation times, diffusion constants, spectral linewidths (full width at half maximum), decay of correlation functions.
- **Units:** [time]$^{-1}$ for rates; [length^2/time] for diffusion.
- **Invariance:** positivity is checked at the level of the reduced dynamics (GKSL) and is independent of coordinates; report temperature or drive dependence as part of the DoV.

ϕ — **phase/order parameter.**

- **Proxies:** interference visibilities, condensate fractions, intensity-phase correlations, order-parameter amplitude from fits.
- **Units:** dimensionless phase or an energy density scale tied to calibration.
- **Invariance:** only *relative* phases and invariant correlators are reported; global phase choices are descriptive, not ontic.

T — **internal clock (differences/rates).**

- **Proxies:** beat frequencies, transition rates, time-difference readouts synchronized by the device; long-time frequencies in the pointer.
- **Units:** [time] for differences, $[\text{time}]^{-1}$ for rates.
- **Invariance:** only differences that survive gauge/coordinate changes are admissible; absolute offsets carry no ontological weight.

χ^μ — **directional order / transport bias.**

- **Proxies:** birefringence-like splittings, anisotropic scattering patterns, direction-dependent mobilities, small-angle transport asymmetries.
- **Units:** device-specific; often dimensionless anisotropy factors or velocities.
- **Invariance:** χ^μ appears through invariant statements about the effective cone g_{eff} (ASCII: `g_eff`); coordinate axes are bookkeeping only.

Criterion 2.1: Proxy Calibration Rule.

A descriptor counts as physically instantiated only if at least one proxy measurement is specified with units, calibration steps, and an explicit domain-of-validity (noise floor, bandwidth, linearity). Proxies must be linked to gauge-invariant readouts (ratios, eigen/spectral data, or invariant correlators).

Checklist

Reporting checklist for symbols ($\rho, \eta, \phi, T, \chi^\mu$).

1. Units and intended scale (with DoV sentence).
2. Proxy device(s) and calibration (time constants, bandwidths).
3. Gauge status (what is invariant; what is a describer choice).
4. Gate tags: S (spectral safety), G (causality within g_{eff}), TQ (complete positivity for reductions), D (non-negative dissipation).

title=ASCII notes for portability

When plain-text is required (e.g., metadata, filenames), we use `rho`, `eta`, `phi`, `T`, and `chi^mu` for χ^μ; the effective cone is `g_eff`. The mathematics in the PDF remains canonical.

2.2 Observables and Units

2.2.1 Finite, Unit-Bearing

The business end of physics is the *pointer*. A model becomes physical only insofar as it can be connected to a procedure that returns numbers with units, within a stated domain of validity (DoV), and whose outcomes are stable under coarse–graining. This section fixes our usage of "observable" and the unit discipline we will enforce throughout the volume.

Definition 2.2: Pointer Observable (operational).

A quantity O is a pointer observable if there exists a conceived measurement pipeline (preparation $\to$ coupling $\to$ readout) that, in principle, yields a *finite* outcome $o \in \mathbb{R}$ with explicit *units*, is gauge/coordinate invariant at the level of the reported descriptor, and whose statistics are *stable* under coarse–graining below device resolution. The pipeline must be compatible with causal open dynamics and the S/G/TQ/D gates.

Criterion 2.2: Finite–Outcome Criterion.

If the formalism predicts $|O| \to \infty$ or undefined units for admissible states within the declared DoV, the claim is *not* about a physical observable. Either the quantity is merely a descriptor, the chart is pathological, or the model has reached end–of–validity and requires an effective replacement.

Criterion 2.3: Unit–Bearing Criterion.

Every reported observable comes with (i) units (SI by default), (ii) the calibration map from raw readout to the quantity with units, and (iii) the device resolution and acquisition timescale that bound the coarse–graining. Natural–units shortcuts (e.g. setting $c = \hbar = 1$) are allowed in the algebra but the *dimensional class* (e.g. energy, time, length) must be stated explicitly when reporting physical claims.

title=From device to number: a minimal template

A typical pipeline reads

$$\text{raw counts } n \xrightarrow{\text{calibration}} r = \alpha n + \beta \xrightarrow{\text{normalization}} O = \mathcal{C}(r;\, \theta)\,,$$

where α, β and θ are calibration parameters determined by reference runs. The reported O carries units via $\mathcal{C}$, e.g. $O = f$ in Hz, $O = \varepsilon$ in J/m^3, or a dimensionless *bounded* ratio such as $p \in [0, 1]$ (energy share, polarization fraction). Stability is checked by repeating the pipeline under coarser binning and longer integration; the result must not drift beyond stated tolerances.

Three classes of observables we will use repeatedly.

(a) *Intensive field–like quantities:* energy density ε (J/m^3), effective stiffness ρ (e.g. Pa), viscosity η (Pa s), diffusion constant D (m^2/s).

(b) *Kinematic/time–scale quantities:* frequency f (Hz), relaxation time τ (s), damping rate γ (1/s). These appear naturally in open–system generators (e.g. GKSL rates) and in our arrow–of–time gate (non–negative dissipation).

(c) *Bounded ratios (dimensionless but physical):* birefringence contrast $\Delta n/n$, polarization fraction Π, and energy shares $p_i = E_i/E_{\text{tot}} \in [0, 1]$. Ratios are acceptable *physical* observables when their numerators and denominators are separately measurable and the ratio remains within finite bounds across the DoV.

Table 2.1: Examples of pointer observables and units (indicative only).

Observable	**Operational proxy (pointer)**	**Units**
Energy density ε	Calorimetric readout per volume; spectral integration	J/m^3
Effective stiffness ρ	Stress–strain slope in linear regime	Pa
Viscosity η	Ring–down width; Stokes drag	Pa s
Relaxation time τ	Exponential fit of envelope $e^{-t/\tau}$	s
Frequency f	Zero–crossings per time; FFT peak	Hz
Directional order χ^μ	Birefringence/anisotropic scattering	dimensionless (bounded)
Energy share p	Channel power over total power	dimensionless in $[0, 1]$

Gate checks for observables. Whenever an observable is introduced, we will attach a brief S/G/TQ/D audit:

- **S (spectral safety):** reported spectra of pointer observables contain no negative or ghost modes.
- **G (causality within g_{eff}):** the measurement protocol respects the effective cone; inferred transport speeds lie within it.
- **TQ (complete positivity):** reduced dynamics for the device plus subsystem is completely positive; no negative probabilities are induced by the coupling.
- **D (non–negative dissipation):** extracted rates γ or diffusion constants are non–negative; the arrow of time is encoded as entropy production or monotone decay.

title=Unit discipline and natural units

Internally we often compute in natural units to streamline algebra, but *claims* are reported with explicit dimensions using `siunitx`. Thus a result derived with $c = \hbar = 1$ will still be stated as, e.g., $f = 5.3\,\text{THz}$ or $\varepsilon = 2.1 \times 10^6\,\text{J/m}^3$. Where only ratios matter, we will say so explicitly and justify the choice by gauge/scale invariance.

Criterion 2.4: Device–Stability Criterion.

Let O_Δ denote the outcome reported at time–bin width Δ (or a spatial voxel of size Δ). There exists Δ_0 such that for all $\Delta \geq \Delta_0$,

$$\frac{|O_\Delta - O_{\Delta_0}|}{\max\{1, |O_{\Delta_0}|\}} < \varepsilon\,,$$

with tolerance ε fixed by calibration. Failure indicates either insufficient coarse–graining or a model outside its DoV.

The next subsection puts these ideas to work in a "math-to-measurement" checklist we will attach to every new quantity introduced in the book.

2.3 Gauge-Invariant Descriptors

2.3.1 Reality via Invariance

Physics does not speak the language of coordinates or gauge choices; it speaks the language of *invariants*. A descriptor earns candidacy for "reality" only insofar as it is insensitive to admissible coordinate/gauge transformations *and* can be bridged to a pointer readout.

Definition 2.3: Gauge-Invariant Descriptor.

Let $\mathcal{S}$ be the state space and $\mathcal{G}$ the (internal and/or spacetime) gauge group acting on $\mathcal{S}$. A map $D : \mathcal{S} \to \mathbb{R}^k$ is a gauge-invariant descriptor if

$$D(g \cdot s) = D(s) \qquad \forall\, s \in \mathcal{S},\ \forall\, g \in \mathcal{G},$$

and there exists a pointer-bridge Π that calibrates D into finite, unit-bearing outcomes within a declared domain of validity.

Principle 2.1: Reality via Invariance.

If two descriptions differ only by an admissible gauge/coordinate transformation, they represent the *same* physical situation. Only quantities that survive such transformations are candidates for ontology; gauge-dependent choices are bookkeeping.

title=Declare the gauge before you claim

For every model claim we will (i) state the relevant group(s) $\mathcal{G}$ and the transformation rules; (ii) identify the combinations that are invariant; (iii) give a pointer-bridge from those invariants to a device readout; (iv) attach S/G/TQ/D checks. This turns "gauge realism" into workflow, not slogan.

Canonical exemplars.

(a) *Electromagnetism.* The potential A_μ changes as $A_\mu \mapsto A_\mu + \partial_\mu \Lambda$, while the field tensor $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is invariant. So are the scalars $F_{\mu\nu}F^{\mu\nu}$ and $\tilde{F}_{\mu\nu}F^{\mu\nu}$. The Aharonov–Bohm phase $W[C] = \exp\Big(iq \oint_C A_\mu dx^\mu\Big)$ is gauge invariant (a Wilson loop) and has a direct pointer-bridge via interferometric fringe shifts.

(b) *General Relativity.* Metric components $g_{\mu\nu}$ depend on charts; curvature scalars such as R, $R_{\mu\nu}R^{\mu\nu}$, and the Kretschmann $K = R_{\mu\nu\alpha\beta}R^{\mu\nu\alpha\beta}$ are invariant. Coordinate singularities (e.g. at an event horizon) disappear in invariants; genuine pathology appears as divergence of these scalars and thus signals *end-of-validity* rather than a "new object".

(c) *Quantum phases and order parameters.* A global phase of a wavefunction is unobservable; gradients of a condensate phase couple via the gauge-covariant derivative $(\nabla - iqA)$ and enter invariant energies like $\int |(\nabla - iqA)\psi|^2\, d^3x$. Topological winding numbers (vortex charge) and Wilson phases anchor the observable content.

(d) *Effective directional order.* Our descriptor χ^μ encodes anisotropy of transport. Its *gauge-free content* is carried by invariants such as $\chi_\mu \chi_\nu\, g_{\text{eff}}^{\mu\nu}$ and by birefringence

ratios $\Delta n/n$ obtained from pointer devices; any particular coordinate representation of χ^μ is not ontic.

Criterion 2.5: Gauge Report (minimal).

Every new descriptor introduced in this volume comes with a *gauge report*: (1) the acting group(s) $\mathcal{G}$; (2) its transformation rules; (3) the invariant combinations we will use; (4) the pointer-bridge and units; (5) S/G/TQ/D checks. Claims missing a gauge report are regarded as descriptive, not physical.

Building invariants in practice. Given primitive variables X with known transformation rules:

i Identify the group $\mathcal{G}$ and covariant objects (tensors, forms, sections of bundles).

ii Contract indices with invariant structures (e.g. $g_{\text{eff}}^{\mu\nu}$, Levi–Civita) to form scalars or gauge-invariant tensors.

iii Prefer *spectral* invariants (eigenvalues of self-adjoint operators, singular values) because they map cleanly to devices.

iv Provide the calibration from these invariants to pointer observables; declare the domain of validity and units.

Table 2.2: Illustrative invariants and their operational proxies.

Invariant descriptor	**Pointer proxy**	**Units**
$F_{\mu\nu}F^{\mu\nu}$ (EM)	Field energy density via calorimetry	$\mathrm{J/m^3}$
Wilson phase $W[C]$	Interference fringe shift	dimensionless (bounded)
Kretschmann K (GR)	Tidal accelerations (geodesic deviation)	$\mathrm{m^{-4}}$ (effective)
$\chi_\mu\chi_\nu\, g_{\text{eff}}^{\mu\nu}$	Birefringence ratio $\Delta n/n$	dimensionless
Spectral gaps $\{\lambda_i\}$	Transition frequencies/lines	Hz

title=Common pitfalls and cures

- *Reifying charts.* Reporting g_{tt} as if it were physical: cure by switching to curvature scalars or redshift factors with a calibration.
- *Reifying potentials.* Attributing reality to A_μ rather than to $F_{\mu\nu}$ or Wilson phases: cure by quoting invariant quantities and their pointer bridges.

- *Hidden gauge in media.* Direction fields or "preferred frames" introduced by hand: cure by stating g_{eff} and reporting invariant contractions and bounded ratios only.

Gate interface. The invariants chosen must still pass the consistency gates: their spectra (when operators are involved) are non-negative for observables (S); their inference respects the effective cone g_{eff} (G); the induced reduced dynamics for device+system is completely positive (TQ); and extracted dissipation rates are non-negative (D). Invariance is necessary, not sufficient.

With these rules, the *reality content* of our formalism is carried by invariant, calibrated descriptors. Coordinates, gauges, and diagrammatic choices remain what they should be: tools of description.

Chapter 3

Method and Style of Claims

This chapter fixes how claims are to be written, checked, and read in this volume. Our goal is to keep abstraction honest and to make every statement actionable: a claim comes with a scope (domain of validity), an account of its idealizations and their stability at finite approximation, and a concrete bridge from model descriptors to pointer observables. Throughout, the S/G/TQ/D gates remain the standing audit.

3.1 Domain of Validity

A result is never true "everywhere"; it is true *within a domain of validity (DoV)* determined by scales, regimes, and approximations. Declaring the DoV is not a formality; it is the first protection against reifying idealizations and the main cue for when an argument must be rephrased.

3.1.1 Declaring DoV

We treat DoV as a first-class citizen: every substantive claim carries a compact DoV sentence and, where appropriate, a quantitative boundary (inequalities on small parameters, bandwidths, intensities, Knudsen-like numbers, etc.).

Definition 3.1: Domain of Validity (DoV)

Let Θ denote controllable knobs (e.g., drive, density, temperature, resolution), $\mathcal{S}$ admissible states, and $\mathcal{A}$ the list of approximations invoked. The DoV of a claim is the subset

$$\mathrm{DoV} = \Big\{(s,\theta) \in \mathcal{S}\times\Theta : \text{ all approximations in } \mathcal{A} \text{ hold to stated tolerances and all gate checks S/G/T}$$

What must be declared.

- **Regime and scales.** The ranges of (s, θ) where the statement is intended (e.g., linear response, weak coupling, paraxial propagation, Markovian bath).
- **Small parameters.** Name the controlling dimensionless quantities and thresholds (e.g., $\epsilon_{\mathrm{nl}} < 0.1$, $\omega/\omega_c < 0.3$, $\mathrm{Kn} < 0.01$).
- **Device envelope.** Noise floor, bandwidth, dynamic range, voxel size or time-bin Δ used to construct pointer observables.
- **Gauge status.** Which combinations are invariant (reported) and which are merely choices (suppressed in the final statement).

Criterion 3.1: DoV Declaration Rule

A claim is *incomplete* unless it includes a one-sentence DoV declaration that states regime, scales, and at least one quantitative boundary (threshold on a controlling parameter or a device envelope).

Template (fill in the brackets).

DoV sentence (author template)

This claim holds for [regime, e.g. linear response, weak anisotropy] *with* [controlling parameters and thresholds, e.g. $\epsilon_{\mathrm{nl}} < 0.1$, $\Delta f < 1\,\mathrm{MHz}$], *on* [scales, e.g. $k\ell < 0.3$, $t > 10\,\mathrm{ns}$], *with pointer data acquired at* [bin width Δ / integration time] *and reported as* [invariant combination]. *Gate checks S/G/TQ/D passed.*

Illustrations.

- **GR (horizon vs. center).** Statements about regularity at an event horizon are made in curvature invariants and hold for geodesic observers with tidal gradients below the device threshold; the DoV excludes neighborhoods where curvature scalars diverge ($r \to 0$).
- **QFT (low-energy effective).** A ϕ^4 effective model is licensed below a cutoff Λ with renormalized couplings defined at scale μ; the DoV requires regulator independence within the quoted uncertainty and excludes Landau-pole neighborhoods.
- **Measurement (linear photodetector).** Linearity holds up to a saturation intensity I_s and within bandwidth B; reported visibilities are invariant ratios and stable under rebinning $\Delta \to 2\Delta$ to tolerance ε.

3.1.2 End-of-Validity

DoV has an edge. When we hit it, we do not "discover" a new entity; we discover the *end of a description.* Recognizing and reporting end-of-validity (EoV) is part of intellectual

honesty and a source of new models.

Principle 3.1: End-of-Validity, not Ontology

A blow-up of a non-invariant quantity, a failure of stability-at-finite for a purported observable, or a violation of S/G/TQ/D is a *domain boundary*, not a new thing in nature. The remedy is reparametrization, coarse-graining, or a new effective theory with an expanded DoV.

Operational triggers for EoV.

T1. **Coordinate artifact.** The divergence disappears in invariants (e.g., g_{tt} blows up, curvature scalars do not).

T2. **Limit misuse.** A result depends materially on the ideal limit; stability-at-finite fails the audit for any reasonable tolerance.

T3. **Gate failure.** Ghost/negative spectra (S); super-cone signalling w.r.t. g_{eff} (G); loss of complete positivity in reduced dynamics (TQ); negative dissipation on the reported timescale (D).

T4. **Parameter exit.** A controlling parameter leaves its threshold band (e.g., $\epsilon_{\text{nl}} \uparrow$ beyond the quoted bound); linear-response assumptions no longer hold.

T5. **Device drift.** Reported values change beyond tolerance under coarser binning or longer integration (violates device-stability criterion).

T6. **Discretization dependence.** Mesh or regulator dependence does not diminish under refinement within the claimed DoV.

Three-step protocol.

i. **Diagnose invariants.** Recast the claim in invariant combinations; if the issue vanishes, annotate it as a chart artifact.

ii. **Audit the limit.** Run the stability-at-finite test with a device-justified tolerance; if it fails, retract the idealization or move the boundary.

iii. **Repair or replace.** If a gate fails or parameters exit, replace the model by an effective one (e.g., add dissipation channels, enlarge state space, change coarse-graining) and declare a new DoV.

Checklist

DoV/EoV reporting checklist.

- One-sentence DoV with regime, scales, thresholds, and device envelope.
- Invariant form of the claim (gauge/coordinate free).

- Stability-at-finite audit if an ideal limit is used.
- Gate check S/G/TQ/D and any failures annotated as EoV.
- If EoV is reached: state which trigger fired and how the model was repaired or replaced.

3.2 Idealization Audits

Idealizations are indispensable: plane waves simplify scattering, delta functions simplify couplings, frictionless limits simplify kinematics, and continuum fields simplify media. But idealizations also tempt reification. Our policy is simple: we will use any idealization as a *tool*, never as an *ontological claim*. The way to keep this discipline is to run an *idealization audit* whenever a proof, computation, or narrative leans on a limit (∞, zero width, perfect symmetry, vanishing noise, exact continuum, infinitesimal step). The audit answers two questions:

1. Does the reported quantity (with units) remain *stable at finite* approximation, within device tolerances?
2. Is the claim *regulator/mesh independent* within the declared domain of validity (DoV), or is the dependence itself part of the reported physics (and then stated as such)?

These questions are methodological, not pedantic. They prevent us from exporting artifacts of charts, meshes, and regulators into ontology, and they localize where a model ceases to be adequate (end-of-validity).

3.2.1 Stability at Finite

We now formalize the stability test. Suppose an idealization introduces a family of *finite* approximants that converge to a formal limit only in a purely mathematical sense. We require that the *reported observable* be insensitive to further tightening of the idealization, beyond a tolerance set by calibration.

Definition 3.2: Finite-Approximation Family

An idealization is accompanied by an indexed family of finite approximants $\{X_\alpha\}_{\alpha\in\mathcal{I}}$ (e.g., packet widths σ, cutoffs Λ, mesh sizes h, time-steps Δt) and a reported observable $O(X_\alpha)$ with units. The index set $\mathcal{I}$ is totally ordered by "finer" (e.g., $\sigma \downarrow 0$, $\Lambda \uparrow \infty$, $h \downarrow 0$).

Relative stability metrics. We will use simple, device-facing metrics; pick one appropriate to the quantity:

$$\delta_{\mathrm{rel}}(\alpha \to \alpha') = \frac{|O(X_{\alpha'}) - O(X_\alpha)|}{\max\{1, |O(X_\alpha)|\}}, \tag{3.1}$$

$$\Delta_{\mathrm{mesh}}(h \to h/2) = \frac{\|O_h - O_{h/2}\|}{\max\{1, \|O_{h/2}\|\}}, \qquad \text{(operator or field norm)}, \tag{3.2}$$

$$\Delta_{\mathrm{spec}}(\Lambda \to \Lambda') = \max_i \frac{|\lambda_i(\Lambda') - \lambda_i(\Lambda)|}{\max\{1, |\lambda_i(\Lambda)|\}}, \ \text{(spectral features)}. \tag{3.3}$$

A tolerance ε is fixed by calibration and the DoV (noise floor, dynamic range, bandwidth). Stability means $\delta_{\mathrm{rel}} < \varepsilon$ (or the corresponding mesh/spectral criterion) once α is "fine enough".

Criterion 3.2: Stability-at-Finite

There exists α_0 such that for all $\alpha' \succ \alpha \succ \alpha_0$, the chosen stability metric (e.g. (3.1)–(3.3)) satisfies $< \varepsilon$ with a tolerance justified by calibration. If not, the idealization is *not* licensed for this claim; the computation should be repeated at finite, or the model/DoV must be revised.

Worked micro-examples (device-facing).

(a) **Plane wave vs packet.** Replace an infinite plane wave by a Gaussian packet of width σ and total energy E; report a detector energy share $p(\sigma)$ in a fixed bandwidth. If $|p(\sigma/2) - p(\sigma)|/\max\{1, |p(\sigma)|\} < \varepsilon$ for $\sigma < \sigma_0$ (device-chosen), the plane-wave idealization is licensed for p.

(b) **UV cutoff in a low-energy model.** Introduce a cutoff Λ and renormalize couplings at a scale μ. Report only those observables whose residual Λ-dependence is below tolerance in the quoted DoV; if not, the observable is outside DoV or the model must be enriched (additional channels).

(c) **Mesh/time-step refinement.** For a PDE observable O_h, compute at $h, h/2$, and (optionally) $h/4$; demand $\Delta_{\mathrm{mesh}}(h \to h/2) < \varepsilon$. Optionally use Richardson extrapolation to estimate the $h \to 0$ value, but report the *finite* result with an error budget.

(d) **Open-system rates under GKSL.** If a predicted rate γ_k flips sign under tighter coarse-graining or longer integration, the D-gate (non-negative dissipation) fails; revise the coarse-graining or model.

Regulator and gauge discipline. The regulator (packet width, cutoff, mesh, step) is a *descriptive* choice, not an ontic parameter. We will:

- Report only invariant combinations (e.g., spectral gaps, energy shares, bounded ratios) that remain finite and unit-bearing.
- Show regulator independence within tolerance; where dependence remains, we either elevate it to *reported physics* (with units and DoV) or retract the claim.

Device-coarsening audit (binning test). Because devices impose finite resolution, we also probe stability under *coarsening* rather than only under mathematical tightening. Let O_Δ denote a reported value at bin width Δ (time or space). Accept if there exists Δ_0 such that for all integers $m \geq 1$:

$$\frac{|O_{m\Delta_0} - O_{\Delta_0}|}{\max\{1, |O_{\Delta_0}|\}} < \varepsilon.$$

This catches claims that are fragile to practical coarse-graining even if they look convergent under ideal refinement.

Error budgets and reporting. Stability statements are meaningless without an error budget. We require: (i) statistical uncertainty (repeatability), (ii) calibration uncertainty (slopes/offsets), (iii) discretization/regularization residual. When possible, we present a small table with these three contributions and the dominant gate-constraint (e.g., D-gate via $\gamma_k \geq 0$).

Audit vs ontology

Passing stability-at-finite *licenses an idealization for this claim in this DoV*. It does not turn the limit object (plane wave, delta, continuum) into a thing in nature. Failing stability does not "discover" a new object; it discovers the end of a description or a missing ingredient in the model.

Minimal audit sheet (attach to figures). For each figure or main numerical claim, we attach:

Idealization Audit (one page)

Idealization: *(packet width / cutoff / mesh / step / symmetry)*
Observable: *units, invariant form, pointer-bridge*
Refinement path: $\alpha \to \alpha'$ *(e.g.,* $\sigma \downarrow$, $\Lambda \uparrow$, $h \downarrow$*)*
Metric & tolerance: *eq.* (3.1) *or* (3.2) / ε = *[value]*
Device coarsening: Δ_0*, bin-multiples tested*
Result: *stability passed/failed; residual dependence quantified*
Gate check: *S (spectra), G (within* g_{eff}*), TQ (CP), D (rates* ≥ 0*)*

DoV note: *scales/regimes; boundary trigger for EoV*

Typical failure modes (and cures).

- **Limit-sensitive observable.** Quantity drifts with regulator beyond tolerance. *Cure:* report a different invariant (e.g., ratio or spectral feature) or expand the model (extra channels).
- **Hidden gauge/coordinate artifact.** Convergence appears only in a particular chart. *Cure:* rewrite in invariant descriptors; if still fragile, mark EoV.
- **Non-CP reduced dynamics.** A mapping inferred from finite data fails complete positivity when extended. *Cure:* return to a legitimate GKSL generator or restrict the map to its DoV.
- **Negative dissipation.** Estimated γ becomes negative at "finer" settings. *Cure:* increase coarse-graining window or include missing dissipation paths.

Gate 3.1: S/G/TQ/D interface for idealizations

S: no ghost modes introduced by the regulator/mesh. **G:** signal speeds and inference remain within g_{eff}. **TQ:** reduced dynamics are CP (rates non-negative; legitimate Lindblad form). **D:** dissipation and diffusion coefficients remain non-negative on the reported timescales.

With these practices, idealizations power the analysis without dictating the ontology. What survives is what is stable, finite, unit-bearing, and invariant—anchored to devices and disciplined by the gates.

3.3 Math-to-Measurement Pipeline

A statement is physical only insofar as it can be carried from formal descriptors to a device readout with units, within a declared DoV, and under the S/G/TQ/D gates. We call this the *math-to-measurement pipeline*. It is not an optional appendix to theory; it *is* the method.

Concretely, a pipeline consists of: (i) model descriptors and their invariants; (ii) an effective dynamics (closed or open) compatible with the DoV; (iii) a measurement coupling and an instrument model; (iv) calibration from raw readout to a unit-bearing observable; and (v) stability and gate checks.

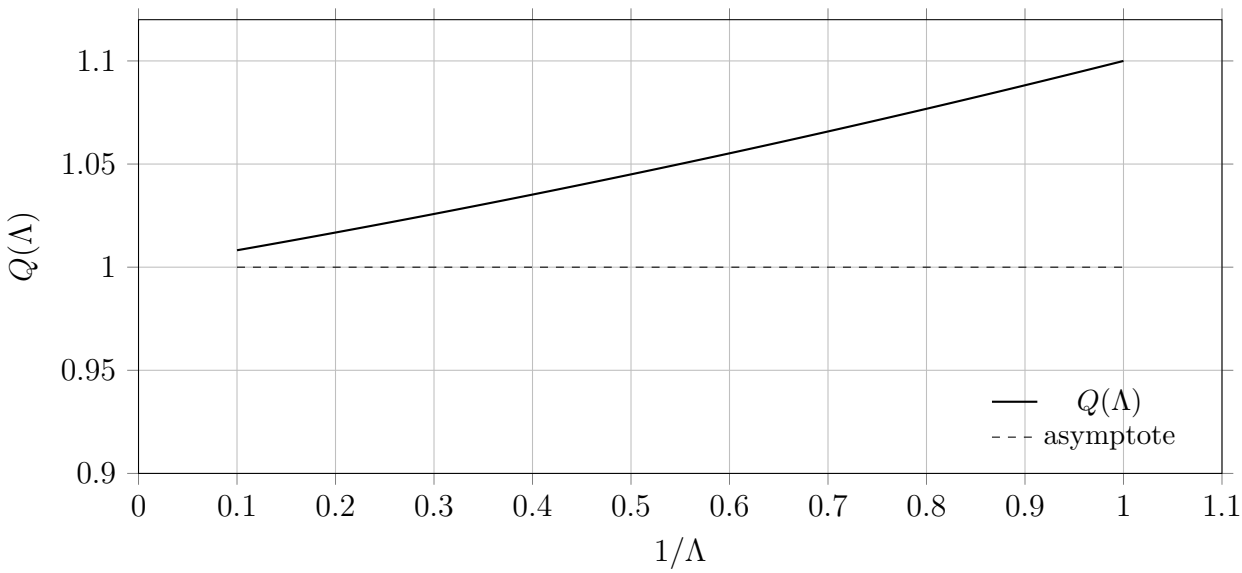

Figure 3.1: Finite-approximation audit: linear/weakly-curved stabilisation in $1/\Lambda$ toward the asymptote (dashed). Declare a tolerance ε and a reporting window in $1/\Lambda$ where the slope/curvature lie below that tolerance.

3.3.1 Model → Pointers

We now fix the objects and maps that make the pipeline explicit.

Definition 3.3: Measurement Pipeline

A measurement pipeline is a tuple

$$\Pi = \big(\mathcal{D}, \mathcal{U}, \mathcal{M}, \mathcal{I}, \mathcal{C}\big)$$

where: (i) $\mathcal{D}$ are *invariant descriptors* with units and DoV; (ii) $\mathcal{U}$ is an *effective dynamics* (closed $-i[H,\cdot]$ or open of GKSL type) on admissible states; (iii) $\mathcal{M}$ is a *measurement coupling* (e.g., POVM or pointer operator); (iv) $\mathcal{I}$ is an *instrument response* (transfer function, bandwidth, noise); and (v) $\mathcal{C}$ is a *calibration map* that turns raw readouts into unit-bearing observables. The reported quantity O is $O = \mathcal{C}\circ\mathcal{I}\circ\mathcal{M}\circ\mathcal{U}$ applied to the descriptors in $\mathcal{D}$, and must pass S/G/TQ/D.

(1) Model descriptors and DoV. State the invariant quantities and the DoV sentence (regime, scales, thresholds, device envelope). Suppress purely descriptive choices (charts, gauges) from the ontology and keep them in the notebook.

(2) Effective dynamics. Closed dynamics: $\dot{\rho} = -i[H(\mathcal{D}), \rho]$. Open dynamics (admissible on subsystems):

$$\dot{\rho} = -i[H(\mathcal{D}), \rho] + \sum_k \gamma_k(\mathcal{D})\Big(L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\}\Big), \qquad \gamma_k \geq 0, \tag{3.4}$$

ensuring complete positivity (TQ) and non-negative rates (D). Causal admissibility (G) is enforced by staying within the effective cone g_{eff}.

(3) Measurement coupling. Specify either a POVM $\{E_j\}$ with probabilities $p_j = \mathrm{Tr}(\rho E_j)$, or a pointer operator O with expectation $\langle O \rangle = \mathrm{Tr}(\rho O)$. If the coupling is weak and continuous, record the integration window and back-action model.

(4) Instrument response. Let K be the (causal) transfer kernel and ν the noise:

$$s(t) = \int_0^\infty K(\tau)\, \langle O(t-\tau) \rangle\, d\tau + \nu(t). \tag{3.5}$$

For counting devices, raw counts are $n = \int_{t_0}^{t_0+T} s(t)\, dt$ with T the integration time; for imaging, convolve with a PSF and voxelize.

(5) Calibration to units. Use a map $O_{\text{phys}} = \mathcal{C}(n; \theta_{\text{cal}})$ that encodes slopes/offsets, responsivity, and normalization to physical units (Hz, J/m^3, etc.). Report the calibration uncertainty and the device resolution.

Criterion 3.3: Pipeline Completeness

A claim is *pipeline-complete* if it presents $(\mathcal{D}, \mathcal{U}, \mathcal{M}, \mathcal{I}, \mathcal{C})$, declares DoV, and passes S/G/TQ/D. Claims missing any component are *descriptive* until repaired.

Timescales and stationarity. List intrinsic timescales (drive, damping γ_k^{-1}, dephasing) and instrument timescales (sampling, integration, K's support). When reporting steady values, state whether they are long-time limits of (3.4) or finite-time averages; mismatch here is a common source of confusion.

Device-facing stability. Repeat the pipeline with coarser binning $T \to mT$ or broader voxels and verify the device-stability criterion (cf. **??**). If the result drifts beyond tolerance, the claim is outside DoV or the instrument model is incomplete.

Worked toy pipeline (two-level, photodetection). *Descriptors.* Transition frequency ω and drive Ω (units rad/s); dissipation rate $\gamma \geq 0$. *Dynamics.* (3.4) with $H = \frac{\omega}{2}\sigma_z + \frac{\Omega}{2}\sigma_x$, $L = \sigma_-$, rate γ. *Coupling.* Counting of emitted quanta $\Rightarrow$ proxy

for excited population. *Instrument.* Finite-band counter with kernel K; raw counts $n = \int_0^T (K * \langle \sigma_+ \sigma_- \rangle)(t)\, dt + n_{\mathrm{bg}}$. *Calibration.* $r = (n - n_{\mathrm{bg}})/T$, $p = r/(\kappa + r)$ with a known gain $\kappa \Rightarrow$ dimensionless energy share $p \in [0, 1]$. *Report.* p is bounded, unit-bearing via calibration, and stable under $T \to 2T$ within ε. *Gates.* S: no negative spectra; G: no super-cone signalling; TQ: Lindblad form with $\gamma \geq 0$; D: non-negative rate.

From math to pointer – minimal checklist

1. Invariant descriptors and one-sentence DoV.
2. Effective dynamics stated with timescales (closed or GKSL).
3. Measurement coupling (POVM/pointer) and integration window.
4. Instrument response K and noise model ν.
5. Calibration map to units and error budget.
6. Stability under coarser binning and regulator refinement.
7. Gate audit S/G/TQ/D and boundary trigger for EoV.

Reporting style (what the reader sees). We will attach to each substantive claim a two-line caption-level statement: the invariant quantity with units and DoV, followed by the pipeline sketch ("H, L_k, γ_k; POVM E_j; K; calibration $\mathcal{C}$") and the gate verdict. This makes claims citable and reproducible.

Gate 3.2: S/G/TQ/D interface for pipelines

S: Observables have non-negative spectra; simulated/fit spectra do not introduce ghosts. **G:** Inference respects g_{eff}; no superluminal channels in $\mathcal{U}$ or $\mathcal{I}$. **TQ:** Maps on subsystems are CP (Lindblad-compatible; rates ≥ 0). **D:** Dissipation/linewidths are non-negative on reported timescales; the arrow of time is encoded as entropy production or monotone decay.

With the pipeline explicit, there is no gap between symbols and devices. Mathematics does not merely decorate the claim; it transports it to the pointer, where physics lives.

Chapter 4

The Spectrum-First Principle

The working stance of this volume can be stated more pointedly: *Spectrum first.* Physical modelling should begin with spectral fields and mode distributions; "particles" and "objects" are special, *stabilized* spectral organizations (localized packets, solitons, poles) within that prior spectral medium. Devices read spectra first (power spectral densities, resonance lines, scattering poles); object-like claims are admitted only when the underlying spectral picture licenses them on a declared DoV.

4.1 Statement and Rationale

Principle 4.1: Spectrum First

Start from *spectral descriptors* (mode families, shares, poles, bandwidths) and their dynamics; treat object-like entities (tracks, lumps, "particles") as *emergent* spectral organizations that pass stability, causality, and positivity checks on the DoV. Devices couple to spectra; objects are what spectra sustain.

title=Why spectrum first? (device, stability, gates)

Device. Pointers are frequency–band instruments by design; even "counts" are filtered spectral shares. **Stability.** Localized objects persist when their spectral composition is robust to coarse-graining and dissipation. **Gates.** S: non-negative spectral densities; G: front/cone constraints on transport; TQ: CP when reducing to sub-spectra; D: non-negative entropy on the window.

4.2 From Spectral Fields to Particles (Solitons)

Let $\{m\}$ index calibrated modes (continuous or discrete). Define the *share vector*

$$p_m(t) \;=\; \frac{\langle E_m(t)\rangle}{\sum_n \langle E_n(t)\rangle}, \qquad p_m \geq 0,\ \sum_m p_m = 1,$$

with E_m the mode energy inferred from the response model (units and calibration supplied).

Definition 4.1: Spectral localization and soliton proxy

A configuration is *object-like* on a window if its share vector is spectrally localized and dynamically stable, e.g. by a large inverse participation ratio $\mathrm{IPR} = \sum_m p_m^2$ and bounded drift of p_m under admissible coarse-graining and dissipation (D). A *soliton* is an attractor of the spectral flow with localized p_m and preserved invariants (e.g. topological index).

Criterion 4.1: Spectral anchoring for object claims

An object claim is admitted if the author supplies: (i) the share vector (or a proxy: peak set, residues, poles); (ii) a stability-at-finite audit of p_m (or of the poles/residues) under window/band variations; (iii) a gate line (S/G/TQ/D); (iv) a pointer bridge (how the shares/poles are read).

Pole picture (unified view). In scattering/response, objects correspond to poles of a calibrated Green function or S-matrix; residues encode couplings (units). S demands positive spectral density; G enforces retarded support; TQ appears when we integrate out subsystems; D constrains linewidths via non-negative absorption/EP.

4.3 Spectral Immersion and the Classical Limit

Immersion. Coupling to many modes spreads phase information; off-diagonal coherences dephase, and pointer summaries reduce to *shares*. This is our device-facing face of "decoherence": probabilities appear as *stable energy shares* (Part 9).

Classicality as spectral dominance. A macroscopic body "looks classical" when a few broad spectral structures dominate shares and remain stable under coarse-graining and dissipation on the instrument's timescale. Interference is suppressed at the pointer because competing shares do not survive the immersion dynamics on the DoV.

4.4 Device Proxies and Falsifiable Signatures

Rule 4.1: Device–level proxies for spectrum-first claims

P1. Response spectra: one-sided PSD $S_{YY}(\omega)$; peak sets, Q factors, residues.

P2. Scattering data: differential cross-sections vs energy; pole extractions.

P3. Time–frequency maps: fronts/isochrons (cone), instantaneous bandwidth.

P4. Open quantum shares: $\mathbf{s}(T)$ and their stabilization to $\mathbf{p}$.

Criterion 4.2: Falsifiers (what would refute spectrum-first)

(a) **Point–particle invariance across scale:** a source whose calibrated form factor remains structureless (no broadening, no run of residues) across orders of magnitude in energy, contrary to spectral localization.

(b) **Macroscopic interference without spectral protection:** stable, large-scale interference fringes when the share vector is demonstrably unstable under immersion on the stated DoV.

(c) **Super-cone transport of spectral weight:** pointer–level no-signalling violated by spectral flow outside the inferred cone.

4.5 Relation to Gates and DoV

title=Gate line for spectrum-first

S (spectral safety): non-negative spectral densities; no ghost modes.
G (causality): fronts and isochrons consistent with g_{eff}; retarded support.
TQ (CP): CP reduction when forming shares over sub-spectra or ancillas.
D (dissipation): linewidths/entropy production non-negative on the window.

DoV declaration. Every spectrum-first claim must name band, resolution, acquisition time, and nonlinearity/memory thresholds; "object" language is confined to windows where spectral localization and stability are demonstrated.

Rule 4.2: Author checklist (spectrum-first edition)

S1. Provide the spectral descriptor (shares, peaks, poles) with units and calibration.

S2. Give a finite-audit plot/table vs window/band; report tolerance.

S3. Attach the gate line S/G/TQ/D and the DoV.

S4. Translate any "object" statement back into spectral terms and name a falsifier.

Bottom line. On *spectrum first*, devices read spectra; "objects" are what spectra *successfully* sustain—localized, stable organizations that pass the gates on a declared domain of validity.

Part II

The Abstract and the Real

This part draws a hard but useful line: mathematics provides *forms*; physics licenses only those forms that can be anchored to measurement. Our aim is not to diminish abstraction—we rely on it—but to give it *operational semantics*. The outcome is a discipline: report *invariant, finite, unit-bearing* quantities with a clear bridge to a pointer and a declared domain of validity (DoV), and keep all other descriptive choices in the notebook.

Chapter 5

Operational Semantics of Physical Subject

A central promise of this volume is clarity about "what counts as a *physical* subject." We do not settle this by metaphysical decree but by *method*: a subject is physical exactly to the extent that it can be embedded in a measurement pipeline (cf. Part 3) and survive gauge changes, coarse–graining, and the S/G/TQ/D gate audit. In this chapter we fix the basic definitions and notational discipline that the rest of the part will use.

5.1 Definitions

Abstraction travels fast; devices move carefully. Our definitions therefore pair every model notion with its measurement reading. We begin by isolating the "admissible measurement procedure" that underwrites physical talk, and then contrast physical with mathematical subjects. Throughout, descriptors (**??**) remain the basic carriers of content; they only earn *reality* through invariance and a pointer bridge.

5.1.1 Physical vs Mathematical

Why this matters. Many crises in interpretation start when tools of thought—limits, exact symmetries, convenient coordinates—are promoted to ontology. We avoid this by insisting that physical talk rides on *procedures* that yield finite, unit-bearing outcomes and are consistent with causal, completely positive open dynamics within the declared DoV.

Definition 5.1: Admissible Measurement Procedure

An admissible measurement procedure consists of:

1. **Preparation** of admissible states within the DoV (regime & scales declared).
2. **Coupling** to a pointer (POVM $\{E_j\}$ or pointer operator O) compatible with the

effective cone g_{eff} (causality gate G).

3. **Dynamics** that is closed ($-i[H,\cdot]$) or open of GKSL type with non-negative rates (gates TQ & D).
4. **Instrument response** (transfer kernel K and noise model ν) with stated bandwidth and resolution.
5. **Calibration map** to a *finite, unit-bearing* outcome and an error budget (statistical, calibration, discretization).

A reported observable is the output of this chain and must pass S/G/TQ/D.

Definition 5.2: Physical Subject (operational)

A subject counts as *physical* if there exists an admissible measurement procedure (**??**) that, in principle, returns a *finite* outcome with *units*, is *gauge/coordinate invariant* at the level of the reported descriptor, and is *stable under coarse–graining* below device resolution within the declared DoV.

Definition 5.3: Mathematical Subject

A subject is *mathematical* if it is defined purely within the formalism (e.g., as a limit object, an exact symmetry idealization, or a coordinate artifact) with no admissible measurement procedure attached. Mathematical subjects guide reasoning and approximation; they are not, by themselves, physical claims.

Operational contrast. Table 5.1 summarizes the workflow difference. The point is not to demote mathematics, but to *license* which parts of it enter physical discourse.

Micro–examples (grounding the contrast).

- *EM potentials vs. fields.* A_μ is descriptive; $F_{\mu\nu}$ and Wilson phases are invariant and anchor interferometric readouts.
- *Plane waves vs. packets.* Infinite plane waves are ideals; finite-energy packets deliver bandwidth-limited energy shares at a detector.
- *Coordinate regularity in GR.* g_{tt} may mislead at a horizon; curvature scalars and tidal readouts survive coordinates.

Table 5.1: Operational contrast between physical and mathematical subjects.

	Physical subject	**Mathematical subject**
Definition	Has an admissible measurement procedure; yields finite, unit-bearing, gauge-invariant outcomes within a DoV.	Defined inside the formalism (limit, ideal, chart); no procedure to units is specified.
Gauge status	Reported in invariants (e.g., spectral data, curvature scalars, bounded ratios).	Often gauge/coordinate dependent by construction (requires translation to invariants).
Limits/ideals	Allowed only if *stable at finite* within tolerance.	May rely on exact limits without a stability audit.
Gate audit	Must pass S/G/TQ/D.	Not applicable (until a pipeline is provided).
Outcome	Pointer observable with units and error budget.	Formal descriptor or construct used for analysis.

Criterion 5.1: Physicality Test (minimal)

A statement is *physical* iff it provides (i) an admissible measurement procedure, (ii) an invariant, unit-bearing outcome with a DoV sentence, and (iii) passes S/G/TQ/D. Otherwise the statement is *descriptive* or *mathematical* until a pipeline is supplied (cf. **????**).

Common pitfalls and cures

- **Reifying ideals.** Treating limits (∞, exact symmetry) as objects. *Cure:* run stability-at-finite; report finite proxies.
- **Chart ontology.** Reading coordinates as reality. *Cure:* switch to invariants and pointer bridges.
- **Gate neglect.** Using successful algebra with ghosts/negative rates. *Cure:* attach S/G/TQ/D audit or retract the claim to math.

How this chapter will be used. Sections to follow apply these semantics to *gauge realism* (invariance tests that license reality claims) and to *bridging maps* (what to do when measurability is delayed). The goal is a repeatable workflow: start in symbols, pass through admissible dynamics and instruments, and arrive at a pointer—or stop and mark the statement as mathematical.

5.2 Gauge Realism

The reality we license in this volume is the reality of *invariants*. Coordinates and gauges are instruments of description; what persists under admissible transformations and reaches a pointer with units earns the name "physical." Gauge realism is therefore not a slogan but a workflow: state the acting group, build invariant descriptors, and give the pointer bridge.

5.2.1 Invariance Tests

We formalize the checks that keep coordinate/gauge choices out of ontology.

Definition 5.4: Gauge Setting and Orbits

Let $\mathcal{S}$ be the state space and $\mathcal{G}$ the admissible gauge/ coordinate group acting on $\mathcal{S}$. The orbit of $s \in \mathcal{S}$ is $[s] = \{g \cdot s : \ g \in \mathcal{G}\}$. A map $D : \mathcal{S} \to \mathbb{R}^k$ is *gauge-invariant* iff $D(g \cdot s) = D(s)$ for all $g \in \mathcal{G}$. A pointer observable is licensed only if it factors through the orbit space $\mathcal{S}/\mathcal{G}$ and carries units within a declared DoV.

Principle 5.1: Reality via Invariance (Operational)

Descriptions that differ by an admissible element of $\mathcal{G}$ represent the same physical situation. Quantities that change under $\mathcal{G}$ are bookkeeping; only invariants are candidates for ontology, and only when bridged to a pointer with units within the DoV.

Four invariance tests (to be applied by default). Let Q be the reported quantity and ε the tolerance justified by calibration.

1. **Chart–swap test.** Re-express the claim in two admissible coordinates (or frames). Accept if the numerical report for Q agrees within ε after proper unit conversion.

2. **Gauge–kick test.** Apply an explicit small gauge transform g_α (phase shift, gauge function Λ, reparametrization). Accept if Q is unchanged to within ε for all admissible α.

3. **Spectral test.** Replace raw components by spectral invariants: eigenvalues of self-adjoint operators, singular values, curvature scalars, Wilson phases. Accept if the report can be fully written in spectral data.

4. **Quotient test.** Exhibit a map $\tilde{Q} : \mathcal{S}/\mathcal{G} \to \mathbb{R}^m$ such that $Q = \tilde{Q} \circ \pi$ where $\pi : \mathcal{S} \to \mathcal{S}/\mathcal{G}$ is the orbit projection. This is the structural way to show gauge–independence.

Criterion 5.2: Minimal Gauge Report

Every new descriptor in this part carries a *gauge report*: (1) acting group(s) $\mathcal{G}$; (2) transformation rules; (3) invariant combinations we will use; (4) pointer bridge and units; (5) S/G/TQ/D checks. Missing items demote the statement to descriptive status.

Canonical exemplars (worked at the pointer level).

- *Electromagnetism.* $A_\mu \mapsto A_\mu + \partial_\mu \Lambda$; invariants: $F_{\mu\nu}$, scalars $F_{\mu\nu}F^{\mu\nu}$ and $\tilde{F}_{\mu\nu}F^{\mu\nu}$, and Wilson phase $W[C] = \exp\Big(iq \oint_C A_\mu dx^\mu\Big)$. Pointer bridges: calorimetric energy density, fringe shifts in interferometry (for $W[C]$). DoV: device bandwidth and linearity.

- *General Relativity.* Components $g_{\mu\nu}$ are chart–dependent; invariants: scalar curvature R, Kretschmann $K = R_{\mu\nu\alpha\beta}R^{\mu\nu\alpha\beta}$, redshift factors along geodesics. Pointer bridges: tidal acceleration patterns (geodesic deviation), frequency shifts. Horizon regularity is stated in invariants; blow-ups of g_{tt} alone are artifacts.

- *Order parameters and phases.* Global phase is descriptive; gauge– covariant gradients $(\nabla - iqA)\psi$ enter invariant energies $\int |(\nabla - iqA)\psi|^2 \, d^3x$. Pointer bridges: interference visibilities, persistent current quantization (winding numbers).

- *Effective directional order (our χ^μ).* Report only invariants such as $\chi_\mu \chi_\nu \, g_{\text{eff}}^{\mu\nu}$ and bounded ratios like birefringence $\Delta n/n$. Pointer bridges: analyzer–angle modulation, anisotropic scattering. Coordinate components of χ^μ are not ontic.

Units and invariance. Invariance is compatible with units: many invariants are intensive (unit-bearing) quantities (energy densities, spectral gaps), or bounded dimensionless ratios (phases, polarizations). We will prefer ratios when they reduce device dependence without losing physical content.

Gauge report — author template

Group $\mathcal{G}$: *(e.g., U(1) phase; diffeomorphisms)*
Rules: *(e.g., $A_\mu \mapsto A_\mu + \partial_\mu \Lambda$)*
Invariants: *(e.g., $F_{\mu\nu}$; K; spectral gaps $\{\lambda_i\}$)*
Pointer bridge: *(device, calibration, units)*
DoV: *(regime, bandwidth, thresholds)*
Gates: *S: spectra ok; G: within g_{eff}; TQ: CP map; D: rates ≥ 0*

Micro–proof sketch (factorization lemma). If Q is invariant under $\mathcal{G}$, then Q is constant on orbits and factors through the quotient: there exists a unique $\tilde{Q}$ with

$Q = \tilde{Q} \circ \pi$. Operationally, this licenses reporting $\tilde{Q}$ (with units) as the physical content; any coordinate/gauge choice used in computing Q is a convenience, not ontology.

Table 5.2: From raw descriptors to invariant reports and pointer proxies.

Raw (descriptive)	**Invariant report**	**Pointer proxy**
A_μ (potential)	$F_{\mu\nu}$, Wilson phase $W[C]$	Calorimetry; fringe shift
$g_{\mu\nu}$ (components)	R, K; redshift factor	Tidal pattern; line shift
ψ (global phase)	Winding number; spectral gaps	Josephson/AB interference
χ^μ (components)	$\chi_\mu \chi_\nu \, g_{\text{eff}}^{\mu\nu}$; $\Delta n/n$	Analyzer modulation

Common failure modes (and cures).

- **Chart ontology.** Reporting g_{tt} as a "physical field." *Cure:* move to curvature scalars or measured redshift.
- **Potential reification.** Treating A_μ as ontic. *Cure:* report $F_{\mu\nu}$ or Wilson phases with calibration.
- **Hidden gauge in numerics.** Results visible only in one gauge or mesh. *Cure:* run the chart–swap and gauge–kick tests; rewrite in spectral invariants.

Rule 5.1:

No figure or main claim passes review without a gauge report, an explicit invariant statement with units, and a pipeline sketch to the pointer. If any gate S/G/TQ/D fails, the claim is tagged "descriptive" or marked as end–of–validity for the present model.

Gate 5.1: S/G/TQ/D interface for invariance tests

S: Spectra of reported operators are non–negative for observables. **G:** Inference and measurement lie within the effective cone g_{eff}. **TQ:** Reduced maps built from the coupling/inference are CP. **D:** Dissipation and diffusion rates are non–negative on the reported timescale.

With these tests, gauge realism becomes a reproducible method: identify $\mathcal{G}$, compress to invariants, carry them to a pointer with units, and audit by the gates. What remains invariant—and only that—earns the reality label in this volume.

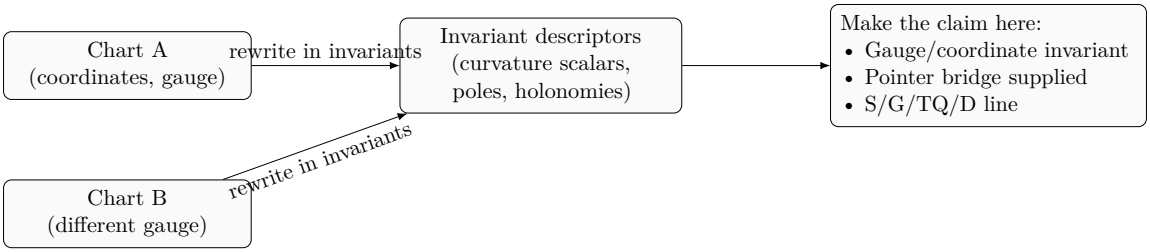

Figure 5.1: Two charts, one physics: claims live in invariants, not in gauge-dependent components.

5.3 Bridging Maps

Between formal descriptors and pointer readouts sits the craft of *bridging*. A bridging map turns invariant, unit-bearing descriptors into predicted device signals under declared conditions (bandwidth, noise, integration windows) and with an error budget. Sometimes the bridge is *direct* (a calibration curve to an ammeter). Sometimes it is *indirect* (an inverse problem via a forward operator). And sometimes it is *delayed*: no device yet exists, but a physically credible pipeline and falsifiable signatures can be specified.

Definition 5.5: Bridging Map

A bridging map is a tuple

$$\mathcal{B} = \big(\mathcal{D},\ \mathcal{U},\ \mathcal{M},\ \mathcal{I},\ \mathcal{C},\ \Sigma\big)$$

with invariant descriptors $\mathcal{D}$ (units & DoV), effective dynamics $\mathcal{U}$ (closed/open, cf. (3.4)), measurement coupling $\mathcal{M}$ (POVM/pointer), instrument response $\mathcal{I}$ (transfer kernel/noise), calibration $\mathcal{C}$ (to units), and a signature bank Σ (predicted patterns: signs/scales/shapes). The reported observable is $O = \mathcal{C} \circ \mathcal{I} \circ \mathcal{M} \circ \mathcal{U}(\mathcal{D})$ and must pass S/G/TQ/D.

Three recurring bridge patterns.

1. **Proxy bridge.** Report a bounded ratio or spectral feature that is already measurable and tightly correlated with the descriptor.

2. **Indirect-inference bridge.** Pose a forward model $y = F(\theta) + \nu$ with θ built from invariants; solve for θ with an identified error budget (Fisher information or posterior width).

3. **Latent-to-pointer bridge.** Evolve under GKSL to long-time pointer frequencies (energy shares) with rates $\gamma_k \geq 0$; use stationarity/ergodicity as the readout principle.

5.3.1 Delayed Measurability

Not every physically meaningful descriptor admits an immediate device. We admit *delayed measurability* claims, but only under strict conditions: the bridge must be specified, signatures must be falsifiable, and a path to proxies or instrument concepts must exist within the DoV.

Definition 5.6: Delayed Measurability

A descriptor D is *delayed-measurable* if there exists a bridging map $\mathcal{B}$ whose $\mathcal{M}, \mathcal{I}, \mathcal{C}$ are specified *conceptually* (not yet realized), together with a set of falsifiable signatures Σ and at least one proxy observable P available now, such that:

signatures in Σ follow from $(\mathcal{D}, \mathcal{U})$ within the DoV, and P correlates with D to declared accuracy.

Principle 5.2: License to Speculate, Obligation to Specify

You may speculate about a descriptor beyond current devices only if you specify (i) a forward pipeline to hypothetical readouts, (ii) falsifiable signatures (sign/scale/-pattern), and (iii) present-time proxies that track the descriptor within a toleranced DoV. Otherwise the claim remains purely mathematical.

Quantifying "how delayed". Bridge Maturity Levels (BML). We classify bridges by *maturity*, analogous to technology readiness:

- **BML-0 (conjectural).** Narrative link only; no proxies; no instrument concept. *Status:* mathematical.
- **BML-1 (proxy-ready).** At least one bounded, unit-bearing proxy P with calibration and DoV; qualitative signature bank Σ.
- **BML-2 (instrument-model).** Forward operator $y = K * O + \nu$ specified; simulated noise, bandwidth, and an error budget; identifiability demonstrated (e.g. non-singular Fisher matrix).
- **BML-3 (lab-path).** Concrete device parameters, calibration plan, and falsifiable thresholds (e.g. $|\Delta| > \varepsilon$) tied to an acquisition time. *Status:* publication-ready prediction.

Criterion 5.3: Acceptance for Delayed Measurability

A delayed-measurability claim is acceptable iff it reaches BML-2 or higher, with (i) a specified forward operator and noise model, (ii) an error budget (statistical, calibration, discretization), (iii) an invariant, unit-bearing proxy today, and (iv)

S/G/TQ/D compliance of the projected pipeline.

Forward operators and identifiability. In indirect bridges we write $y = F(\theta) + \nu$ where θ are invariant, unit-bearing parameters. Local identifiability is certified if the Fisher information $I(\theta) = \mathbb{E}[(\partial_\theta \log p(y \mid \theta))^2]$ is well-conditioned in the DoV, or by a posterior width $\mathrm{Var}(\theta \mid y)$ below a stated tolerance.

Example (proxy first, device later). Directional order χ^μ:

- *Proxy (BML-1).* Birefringence ratio $\Delta n/n$ with analyzer modulation; units: dimensionless; DoV: linear optics, bandwidth B.
- *Instrument model (BML-2).* Time-resolved pump–probe with kernel $K(\tau)$ and noise ν; forward $y = \int K(\tau)\,(\chi_\mu \chi_\nu\, g_{\mathrm{eff}}^{\mu\nu})(t-\tau)\,d\tau + \nu$; demonstrate stability-at-finite under coarser binning.
- *Lab path (BML-3).* Minimum detectable anisotropy $|\Delta n/n| > \varepsilon$ at integration time T; calibration via known retarders; gate audit: spectra non-negative (S), no super-cone signalling (G), CP reduced dynamics (TQ), non-negative rates (D).

When measurability must be delayed. Sometimes the bridge itself demands new scales (e.g., ultra-low noise or high field strengths). We then:

1. Report *only* invariant, bounded signatures Σ (e.g., sign of a shift, parity of a pattern, presence/absence of a gap).
2. Provide a *proxy ladder*: $P_1, P_2, \ldots$ with increasing selectivity and stated correlations to D.
3. Archive the forward model and calibration scripts so that future devices can adopt them unchanged.

Bridge dossier — author template

Descriptor(s): invariants with units & DoV
Bridge pattern: proxy / indirect / latent-to-pointer
Forward operator: $y = K * O + \nu$ (or $F(\theta)$) with device envelope
Signatures Σ: sign/scale/pattern with thresholds
Proxies now: name, calibration, correlation to descriptor
Maturity: BML-0/1/2/3 (justify)
Gates: S/G/TQ/D checks and boundary trigger for EoV

Common pitfalls (and cures).

- **Reifying the bridge.** Treating a forward model parameter as ontic. *Cure:* keep ontology at invariant descriptors; the operator K lives in the instrument model.
- **No error budget.** Predictions quoted without uncertainty. *Cure:* attach statistical/-calibration/discretization components.
- **Gauge leakage.** Using chart-dependent components as "signatures". *Cure:* compress to invariant ratios or spectral data.
- **Gate neglect.** Instrument or coupling induces negative rates or super-cone inference. *Cure:* redesign $\mathcal{M}, \mathcal{I}$ or move the DoV boundary.

Gate 5.2: S/G/TQ/D interface for delayed measurability

S: Reported spectra/ratios are bounded and ghost-free. **G:** Forward operator and inference respect the effective cone g_{eff}. **TQ:** Projected reduced maps are completely positive (Lindblad-compatible). **D:** Predicted dissipation/linewidths are non-negative on stated timescales.

With bridging maps made explicit—and with a discipline for delayed measurability—the theory keeps its promises: every descriptor either reaches a pointer with units (now or by a specified path), or it stays in the mathematics box until it does.

Part III

Infinities, Singularities, End-of-Validity

This part enforces a simple discipline with large consequences: *infinity is a limit device, not a measured value*, singularities are *end-of-validity flags*, and models must say where they hold. We keep the calculational power of asymptotics and idealizations while reporting only *finite, unit-bearing, gauge-invariant* outcomes that pass stability-at-finite tests and the S/G/TQ/D gates.

Chapter 6

Infinity as a Limit, Not a Value

Infinity appears everywhere—in plane waves, deltas, thermodynamic limits, ultraviolet cutoffs. In this volume we adopt the policy already stated in **??**: use ∞ as a *tool* to control approximations; never treat it as an ontic value. The practical question then is: how do we license an ideal limit for *physical reporting*? This chapter answers by specifying legitimate uses, stability checks, and short case studies.

6.1 Legitimate Uses of ∞

Asymptotics let us compress complicated dependences into controlled approximations. But asymptotic symbols ($\to \infty$, $O(\cdot)$) are not self-licensing for physics; they must be tied to *finite* approximants and a device envelope.

Definition 6.1: Operational Asymptotic Statement

Let $\{X_\alpha\}_{\alpha\in\mathcal{I}}$ be a family of *finite* approximants indexed by a refinement parameter (e.g. packet width $\sigma \downarrow 0$, cutoff $\Lambda \uparrow \infty$, mesh size $h \downarrow 0$). A model predicts an *operational asymptotic* for an observable O if there exists a tolerance ε (set by calibration and DoV) and α_0 such that for all $\alpha' \succ \alpha \succ \alpha_0$,

$$\frac{|O(X_{\alpha'}) - O(X_\alpha)|}{\max\{1, |O(X_\alpha)|\}} < \varepsilon.$$

In that case, we may *license* the ideal shorthand (e.g. "$\sigma \to 0$") for reporting O within the stated DoV.

Principle 6.1: Infinity-as-Limit (Operational Reading)

Use ∞ only as a limiter on a family of *finite* constructs for which the reported observable is *stable at finite* within a declared DoV. No observable is ever reported

as ∞; when a formalism produces such a value, it indicates end-of-validity or a non-observable quantity.

Legitimate patterns (with reporting notes).

- **Asymptotic expansions.** Report the *remainder bound* in units: $O(X_\alpha) = O_0 + \delta_\alpha$, with $|\delta_\alpha| \leq B(\alpha)$ and $B(\alpha \succ \alpha_0) < \varepsilon \max\{1, |O_0|\}$.
- **Distributional ideals.** Dirac deltas and plane waves are licensed as *limits of finite packets* (test functions with device bandwidth). Only packet-based, unit-bearing readouts are reported.
- **Thermodynamic/continuum limits.** License only if the intensive observable converges under *finite* volume/mesh growth within tolerance, and the device-scale (voxel, integration window) is declared.
- **Ultraviolet regularization.** Introduce a cutoff Λ or smoother; report observables whose residual regulator dependence is below tolerance on the quoted scales. Otherwise, expand the model or move the DoV.

Criterion 6.1: Reportable–at–Limit

A claim that uses a limit is reportable only if it provides: (i) the finite approximants, (ii) a stability-at-finite audit with tolerance ε, (iii) units and a DoV sentence, and (iv) S/G/TQ/D compliance. Otherwise the limit remains mathematical shorthand and the claim is *not* physical.

6.1.1 Asymptotics, Packets

Plane waves, infinite coherence times, and delta couplings buy algebraic clarity by *sacrificing* finiteness. Our policy is to *replace* them by calibrated *packets* and check that the reported observable is insensitive to further narrowing/widening within the device envelope.

Wave and particle packets (recipe). Let a plane wave idealization be replaced by a family of packets ψ_σ with width σ and total energy E. Let $O(\sigma)$ be a detector-level observable (e.g. energy share in a bandwidth, arrival-time fraction).

1. Fix the device limits: bandwidth B, integration time T, dynamic range, and noise floor (DoV sentence).

2. Choose σ_0 such that for all $\sigma < \sigma_0$,

$$\frac{|O(\sigma/2) - O(\sigma)|}{\max\{1, |O(\sigma)|\}} < \varepsilon.$$

3. Report O at a *finite* $\sigma \leq \sigma_0$ with units and error budget; "plane wave" may appear only as a licensed shorthand.

Delta couplings (recipe). Replace δ by a narrow, normalized kernel κ_ℓ of width ℓ (Gaussian/Lorentzian/compact support) and compute the observable $O(\ell)$. License the contact-limit if $|O(\ell/2) - O(\ell)|/\max\{1, |O(\ell)|\} < \varepsilon$ beyond ℓ_0 fixed by calibration (voxel size, detector psf).

Author template — How to report an asymptotic claim

Idealization: plane wave / delta / thermodynamic / continuum.
Finite family: ψ_σ / cutoff Λ / mesh h / volume V.
Observable & units: O with device envelope (bandwidth B, time T).
Stability check: metric and tolerance ε; value of σ_0, Λ_0, or h_0.
Result: finite reported value with error budget; S/G/TQ/D audit.
Note: "$\to \infty$" or "$\to 0$" appears only as licensed shorthand.

Why packets help philosophy as well as physics. Packets make explicit the *device scale* (coarse–graining), preserve units, and expose where the model ceases to be adequate. They also immunize against reifying mathematical artifacts: it is impossible to "measure" a plane wave or a delta; one measures packet statistics within a DoV.

The next section formalizes the *stability-at-finite* protocols and gives compact tests (mesh, spectral, and device-coarsening audits) that every idealization in this part must pass before it is used in reporting.

title=Summary — Three Roles of Infinity (and What to Do)

(1) Useful idealization (licensed as a limit). Use ∞ (or 0) only as a limiting device *after* demonstrating stability at finite approximants within the device envelope. Report the finite, unit-bearing value; the ideal is a shorthand.

(2) Pathology indicator (end-of-validity). If a purported observable diverges, or stability-at-finite fails for any reasonable tolerance, treat this as *end of validity*, not ontology. Repair by reparametrization, coarse-graining, or an improved effective model (**??**).

(3) Coordinate/gauge artifact. Divergences tied to coordinates or gauge choices

disappear in invariants. Rewrite the claim in gauge/coordinate-invariant descriptors and re-audit (**??**).

Non-reification rule. Never treat ∞ as a measured value or a "thing." Infinity is a tool of analysis; reality is what reaches a pointer as a finite, unit-bearing quantity.

6.2 Stability-at-Finite Tests

The power of idealizations comes from what *survives* when we return to finite descriptions. A claim that leans on a limit ($\to \infty$ or $\to 0$) must therefore demonstrate *stability at finite* for the *reported, unit-bearing observable* within the declared DoV (device envelope, regimes, timescales). This section turns that policy into concrete tests and repeatable protocols.

Tolerance budget and acceptance bands

Let $\varepsilon_{\text{stat}}$ (repeatability), ε_{cal} (calibration), and $\varepsilon_{\text{disc}}$ (discretization/regulator) be fractional contributions to uncertainty for the reported observable O. We set a working tolerance

$$\varepsilon \;=\; \sqrt{\varepsilon_{\text{stat}}^2 + \varepsilon_{\text{cal}}^2 + \varepsilon_{\text{disc}}^2}\,,$$

and use two acceptance bands: *warn* if the stability metric is in $[\frac{1}{2}\varepsilon, \varepsilon)$, *fail* if it exceeds ε. Claims reported in the warn band must include a mitigation note (longer integration, refined calibration, or expanded DoV).

Principle 6.2: Report only what is stable at finite.

Whenever a limit is used in analysis, the *reported* quantity must be computed at a finite regulator/mesh/packet and shown to vary less than ε under further tightening *and* under device-level coarsening consistent with the DoV.

Criterion 6.2: Stability metrics (recap).

We reuse the relative-change metrics introduced earlier:

$$\delta_{\text{rel}}(\alpha \to \alpha') = \frac{|O(X_{\alpha'}) - O(X_\alpha)|}{\max\{1, |O(X_\alpha)|\}}, \quad \Delta_{\text{mesh}}(h \to h/2) = \frac{\|O_h - O_{h/2}\|}{\max\{1, \|O_{h/2}\|\}},$$

and a spectral-window metric

$$\Delta_{\text{spec}}(\Lambda \to \Lambda') = \max_i \frac{|\lambda_i(\Lambda') - \lambda_i(\Lambda)|}{\max\{1, |\lambda_i(\Lambda)|\|}\,.$$

Pass if the relevant metric $< \varepsilon$ beyond some finite threshold α_0, h_0, Λ_0 justified by calibration and the DoV.

6.2.1 Protocols

We collect practical protocols. Each protocol ends with a *finite* report and a short stability log.

P1 — Packet replacement for plane waves and deltas.

1. Replace the ideal by a finite family: plane wave $\Rightarrow$ packets ψ_σ (width σ); delta coupling $\Rightarrow$ normalized kernel κ_ℓ (width ℓ).
2. Fix the device envelope (bandwidth B, integration T, dynamic range, voxel size). Declare the DoV sentence.
3. Compute the unit-bearing observable $O(\sigma)$ (or $O(\ell)$) at $\sigma, \sigma/2$ (or $\ell, \ell/2$); evaluate the relative change δ_{rel}.
4. Choose σ_0 (or ℓ_0) s.t. for all $\sigma < \sigma_0$ (or $\ell < \ell_0$), $\delta_{\text{rel}} < \varepsilon$.
5. **Report** $O(\sigma^\star)$ with $\sigma^\star \leq \sigma_0$ (or $O(\ell^\star)$) and attach the stability log and error budget.

P2 — UV regulator / cutoff independence (low-energy models).

1. Introduce a cutoff Λ (or smoother) and renormalize couplings at scale μ within the DoV.
2. Evaluate the observable at Λ and $\Lambda' = \rho\Lambda$ (e.g. $\rho = 2$): compute Δ_{spec} or δ_{rel} as appropriate.
3. Finite-difference "beta" slope:

$$\beta_O(\Lambda) \equiv \frac{O(\rho\Lambda) - O(\Lambda)}{\log(\rho)} .$$

   Require $|\beta_O(\Lambda)| < \varepsilon \max\{1, |O|\}$ on the quoted scales. If not, the observable is either outside the DoV or the model needs additional channels.
4. **Report** the finite-Λ value with the residual cutoff dependence explicitly quoted.

P3 — Mesh/time-step refinement (continuum/num. PDE).

1. Compute at h, $h/2$, $h/4$ (and Δt, $\Delta t/2$, $\Delta t/4$ if time-stepped). Record $O_h, O_{h/2}, O_{h/4}$.
2. Check decrease factor

$$r = \frac{\|O_h - O_{h/2}\|}{\|O_{h/2} - O_{h/4}\|} \approx 2^p,$$

where p is the formal order. If r deviates strongly from 2^p, refine the scheme or lower the DoV.

3. Require $\Delta_{\text{mesh}}(h \to h/2) < \varepsilon$.
4. **Report** $O_{h/2}$ (or extrapolated value) with a discretization uncertainty derived from the refinement.

P4 — Device coarsening (binning/integration sweeps).

1. With the instrument model fixed, acquire at bin width Δ_0 (or integration time T_0) and at multiples $m\Delta_0$ (or mT_0) for $m = 2, 4$.
2. Check the device-stability criterion:

$$\frac{|O_{m\Delta_0} - O_{\Delta_0}|}{\max\{1, |O_{\Delta_0}|\}} < \varepsilon \quad \text{for all tested } m.$$

3. **Report** the coarsening-insensitive value and note any drift.

P5 — Spectral-window and analysis-choice neutrality.

1. Vary analysis windows (FFT resolution, smoothing bandwidths, filter orders) within the device-allowed range.
2. Demand invariance of reported *spectral features* (peak positions, gaps, bounded ratios) via $\Delta_{\text{spec}} < \varepsilon$.
3. **Report** window/analysis settings and the residual variation.

P6 — Natural-units restoration and dimensional sanity.

1. If algebra used $c = \hbar = 1$, restore SI units *before* stability checks to avoid scale confusion in ε.
2. Confirm that the stability metrics are dimensionless and that device tolerances are expressed in the same units as the reported observable.

Author template — Stability log (attach to figures/claims)

Idealization: packet / cutoff / mesh / time-step / window
Finite settings: $\sigma, \sigma/2$ or $\Lambda, 2\Lambda$ or $h, h/2, h/4$
Device envelope: bandwidth B, integration T, bin Δ
Metric & tolerance: name, value, ε (with budget)
Outcome: pass/warn/fail; finite reported value with units
Notes: S/G/TQ/D audit; DoV boundary or repair if failed

Worked micro-cases.

- **Packet license (plane-wave shorthand).** Energy share $p(\sigma)$ in a fixed bandwidth is computed at σ and $\sigma/2$; relative change < 0.01 with a tolerance $\varepsilon = 0.02$. *License* "plane-wave" shorthand for p in this DoV; report $p(\sigma^\star)$ with units and error.
- **Cutoff residual (low-energy observable).** Observable $O(\Lambda)$ changes by 0.5% when $\Lambda \to 2\Lambda$; with $\varepsilon = 1\%$ the result *passes*; quote the 0.5% residual as a systematic in the error budget.

Criterion 6.3: Fail-fast triggers.

Declare end-of-validity (for this model/DoV) and repair/replace if any holds: (i) stability metric $> \varepsilon$ even at finest feasible settings; (ii) negative dissipation or non-CP reduced dynamics inferred from the data; (iii) super-cone signalling relative to g_{eff}; (iv) strong dependence on analysis windows that cannot be reduced within device constraints.

With these protocols, every use of a limit is disciplined by finite, unit-bearing reports that withstand both mathematical refinement and device-level coarsening. What remains is licensed for physics; the rest either signals end-of-validity or awaits a better model or instrument.

6.3 Case Studies

We close this chapter with compact case studies showing how the *Infinity-as-Limit* policy and the *stability-at-finite* discipline operate in practice. The emphasis is on end-to-end reporting: DoV sentences, finite settings, stability logs, and S/G/TQ/D audits.

6.3.1 UV, Plane Waves

We consider two paradigmatic situations where ∞ often sneaks in: (i) ultraviolet behavior in low-energy effective models, and (ii) the use of plane waves in scattering/propagation. In both, the rule is the same: replace the ideal by a calibrated finite family, show stability within tolerance, and report a *finite, unit-bearing*, gauge-invariant observable.

A. Ultraviolet (UV) in a low-energy effective model. Consider a scalar effective field with Lagrangian density $\mathcal{L} = \frac{1}{2}(\partial\phi)^2 - \frac{1}{2}m^2\phi^2 - \frac{\lambda}{4!}\phi^4$ intended to model dynamics below a physical scale Λ_{phys}. We care about a *finite, unit-bearing* observable, e.g. a spectral gap (or a differential cross section) on kinematic scales $E \ll \Lambda_{\text{phys}}$.

Regulator and renormalization. Introduce a UV regulator $R(p/\Lambda)$ (sharp or smooth) and define renormalized couplings at a scale μ inside the DoV. The reported observable O

may be written as $O(\mu;\Lambda)$ at intermediate steps, but the physical report must exhibit *cutoff insensitivity* within tolerance.

Cutoff-independence audit. Evaluate $O(\mu;\Lambda)$ and $O(\mu;\rho\Lambda)$ (e.g. $\rho = 2$) on the same kinematic point and compute the finite-difference slope

$$\beta_O(\Lambda) \;\equiv\; \frac{O(\mu;\rho\Lambda) - O(\mu;\Lambda)}{\log\rho}\,.$$

With a device-justified tolerance ε, accept if $|\beta_O(\Lambda)| < \varepsilon\,\max\{1,|O|\}$ for $\Lambda \geq \Lambda_0$ (chosen so that spurious regulator artifacts are below noise/calibration). If not, either (i) the observable is outside the quoted DoV, or (ii) the effective model is missing channels and must be extended (new operators, running of parameters, or a lower Λ_0).

Units and invariants. Report invariant, unit-bearing quantities (e.g. pole mass, spectral gap, on-shell cross section with units m^2). Gauge or coordinate choices enter the notebook, not the ontology.

Gate audit. S: spectra non-negative for observables; G: inference within the effective cone g_{eff}; TQ: positivity via an admissible reduced dynamics when open channels are modeled; D: non-negative dissipation (rates ≥ 0) on the reported timescale.

UV case — reporting snippet

DoV: $E \leq 0.2\,\Lambda_{\mathrm{phys}}$, bandwidth B, integration time T.
Finite settings: Λ_0, $\rho\Lambda_0$; renormalization scale μ.
Stability: $|\beta_O(\Lambda_0)| < \varepsilon$ (value and budget stated).
Report: $O(\mu;\Lambda_0)$ with units and residual cutoff systematic.
Gates: S/G/TQ/D passed (notes if any are near-threshold).

B. Plane waves replaced by packets at the detector. Plane waves simplify algebra but carry infinite extent and undefined energy density. We replace them by a family of finite-energy packets $\{\psi_\sigma\}$ of width σ (temporal or spatial) and use a *detector-level* observable as the report.

Observable. Let $p(\sigma)$ be the energy share (dimensionless, bounded in $[0,1]$) delivered into a fixed detector bandwidth B over an integration time T. The calibration maps counts to physical units and then to a bounded ratio.

Stability-at-finite test. Find σ_0 such that for all $\sigma < \sigma_0$,

$$\frac{|p(\sigma/2) - p(\sigma)|}{\max\{1,|p(\sigma)|\}} < \varepsilon\,,$$

with ε set by the noise floor and calibration. Report the finite value $p(\sigma^\star)$ with $\sigma^\star \leq \sigma_0$ and attach the stability log. The phrase "plane wave" is then a *licensed shorthand* for this observable in the stated DoV.

Pipeline consistency. Descriptors $\to$ effective dynamics (possibly open, with non-negative rates in a GKSL generator) $\to$ measurement coupling (POVM/pointer) $\to$ instrument response (kernel K, noise ν) $\to$ calibration to p. Each step is declared with units and timescales.

Plane-wave case — reporting snippet

DoV: linear response, detector bandwidth B, integration T, analyzer setting.
Finite family: packet widths σ, $\sigma/2$ (and optionally $\sigma/4$).
Stability: relative change $< \varepsilon$; device coarsening test passed.
Report: $p(\sigma^\star) \in [0,1]$ with error budget; "plane wave" licensed.
Gates: S (bounded spectrum), G (within g_{eff}), TQ (CP map), D ($\gamma \geq 0$).

C. Failure modes and repairs (UV & packets).

- **UV drift persists.** $|\beta_O| > \varepsilon$ even at large Λ: expand the effective theory (add operators/channels) or shrink the DoV; otherwise mark end-of-validity for this observable.
- **Packet fragility.** $p(\sigma)$ drifts under $\sigma \to \sigma/2$ or device coarsening $T \to 2T$: increase integration time, adjust the bandwidth B, or choose a different invariant observable (e.g. a spectral gap rather than an absolute intensity).
- **Gate violations.** Negative rates in the inferred reduced map, super-cone signalling, or ghost features: revise the coarse-graining, instrument model, or move the DoV boundary.

Table 6.1: UV vs. plane-wave idealizations — diagnosis and repair.

Symptom	**Diagnosis**	**Repair/Report**
O changes with Λ	Outside DoV or missing channels	Extend model or shrink DoV; report residual as systematic
$p(\sigma)$ drifts with σ	Ideal limit not licensed	Increase T/B; choose different invariant observable
Gate S/TQ/D fails	Non-CP map or negative dissipation	Rebuild dynamics/instrument; move DoV boundary
Coordinate artifacts	Chart/gauge dependence	Rewrite in invariants; re-run stability tests

Takeaway. In both UV control and wave propagation, infinity is a *tool* that licenses compact analysis only after finite, unit-bearing reports prove insensitivity within the device envelope. Where insensitivity fails, we do not discover a new "infinite object"; we discover a boundary—*and we say so*.

Chapter 7

Anti-Singularity

Singular behavior in formulas is cheap; singular behavior in nature is rare. This chapter enforces the working policy already stated (**??**): *a divergence is either a coordinate/gauge artifact to be removed, or an* end-of-validity *flag for a model to be replaced—never a new object by itself.* The practical task is to tell the two apart and to report only *invariant, finite, unit-bearing* content with a clear pointer bridge.

7.1 Coordinate vs Curvature

A component that blows up in one chart may be perfectly regular in another. Curvature, by contrast, is chart-free. The first duty, therefore, is to move from components to *invariants* and from narrative to *pointer* quantities (tidal accelerations, redshifts, spectral gaps).

Definition 7.1: Coordinate (removable) singularity

A point or hypersurface is a *coordinate singularity* if there exists a smooth change of variables on an open neighborhood that renders the metric components finite and keeps all scalar curvature invariants and pointer proxies *bounded* within the declared DoV. Typical symptoms: blow-up of specific components (g_{tt}, a lapse, a Jacobian) while curvature scalars remain finite; geodesics cross the surface with well-defined tangent.

Definition 7.2: Curvature (invariant) singularity

A point or region is a *curvature singularity* if along some causal or spacelike approach, at least one scalar polynomial invariant built from the Riemann tensor (e.g. R, $R_{\mu\nu}R^{\mu\nu}$, Kretschmann $K = R_{\mu\nu\alpha\beta}R^{\mu\nu\alpha\beta}$) *diverges*, or if tidal-pointer quantities diverge (unbounded geodesic deviation). Geodesic incompleteness with divergent

invariants signals *end-of-validity* for the current effective theory.

Criterion 7.1: Invariant sanity test

Let $(\mathcal{U}, x^\mu)$ be a chart in which a component diverges. If there exists a smooth recharting $(\tilde{\mathcal{U}}, \tilde{x}^\mu)$ covering the same neighborhood such that all scalar invariants remain bounded and the reported pointer observables stay within device dynamic range, then the divergence is *coordinate* and must be removed from ontology. If no such recharting exists and invariant scalars/ptr-proxies blow up, the model has reached *end-of-validity*.

Geodesic deviation as a pointer bridge. Curvature becomes operational via the relative acceleration of nearby geodesics:

$$\frac{D^2\xi^\mu}{d\tau^2} = -\, R^\mu{}_{\nu\alpha\beta}\, u^\nu u^\alpha\, \xi^\beta, \tag{7.1}$$

with separation ξ^μ and four-velocity u^μ. Accelerometer arrays measure the components on the right-hand side as *finite, unit-bearing* tidal signals; their divergence is physical only if the measured quantities leave device range (DoV declared).

Canonical cases (diagnosis in invariants).

- **Schwarzschild horizon ($r = 2M$).** In Schwarzschild coordinates $g_{tt} \to 0$ and $g_{rr} \to \infty$, yet curvature scalars (e.g. K) are finite; Eddington–Finkelstein/Kruskal charts remove the blow-up. *Verdict:* coordinate singularity; report redshift/tidal invariants instead.

- **Schwarzschild center ($r = 0$).** Kretschmann $K \propto r^{-6}$ diverges; tidal accelerations become unbounded. *Verdict:* curvature singularity $\Rightarrow$ end-of-validity for classical GR; any ontology must await an effective replacement (new DoV).

- **FLRW big bang.** As $a(t) \downarrow 0$, energy density and curvature scalars diverge in classical solutions. *Verdict:* curvature singularity; classical description ends, motivating a new effective regime.

- **Conical defect (cosmic string).** Away from the core, polynomial invariants vanish; globally there is a deficit angle measurable by holonomy (parallel transport). *Verdict:* not a coordinate artifact; curvature is distributional at the core; report the invariant deficit (pointer: lensing).

Chart audit for a suspected singularity

Step 1 — Invariants. Compute R, $R_{\mu\nu}R^{\mu\nu}$, K, and any spectral invariant available; check finiteness within DoV.
Step 2 — Rechart. Move to a regular chart (e.g. EF/Kruskal, Rindler) covering the neighborhood; re-evaluate.
Step 3 — Pointer. Map to tidal/redshift observables and verify device stability (dynamic range, integration window).
Decision. Bounded invariants & stable pointers $\Rightarrow$ coordinate artifact (remove). Divergent invariants/pointers $\Rightarrow$ end-of-validity.

7.1.1 Invariants

We collect the invariant quantities used to ground claims and to separate removable artifacts from genuine pathology.

Definition 7.3: Scalar curvature invariants

Let $R^{\mu}{}_{\nu\alpha\beta}$ be the Riemann tensor. Typical invariants:

$$R, \qquad R_{\mu\nu}R^{\mu\nu}, \qquad K \equiv R_{\mu\nu\alpha\beta}R^{\mu\nu\alpha\beta}, \qquad C_{\mu\nu\alpha\beta}C^{\mu\nu\alpha\beta},$$

with C the Weyl tensor. These are chart-free and serve as diagnostics for singular behavior.

Definition 7.4: Operational tidal invariants

Given a congruence with unit timelike u^{μ}, the *electric* part of the Riemann $E_{\mu\nu} = R_{\mu\alpha\nu\beta}\, u^{\alpha}u^{\beta}$ encodes tidal accelerations accessible to differential accelerometers; its eigenvalues are spectral invariants with *units* (e.g. $1/\mathrm{m}^2$). These enter the pointer bridge via (7.1).

Criterion 7.2: Invariant reporting rule

Report singularity claims in scalar or spectral invariants (e.g. K, eigenvalues of $E_{\mu\nu}$) with units and a device-facing DoV sentence (dynamic range, integration time). Component blow-ups without invariant divergence are *not* admissible as physical claims.

Media and effective cones

In effective media we replace g by an emergent g_{eff}. The same discipline applies: diagnose singular behavior in invariants built with g_{eff} (e.g. contractions with $g_{\text{eff}}^{\mu\nu}$) and report only pointer-anchored quantities (birefringence ratios, causal cones) that remain finite within DoV. Coordinate choices inside the medium are descriptive, not ontic.

With these tools, "singularity" becomes a controlled term of art: either a removable chart feature or a boundary where the current effective description stops being adequate. In both cases, ontology follows invariants—and devices.

7.2 End-of-Validity Flags

A "singularity" in the algebra is not a discovery; it is a *diagnostic*. When a model reaches the boundary of its domain of validity (DoV), it leaves behind recognizable *flags*. The job is to detect these flags in *invariants and pointer observables*, then either rephrase the description (if the issue is merely coordinate/gauge) or replace the model with a better effective theory. This section collects the flags; the next subsection gives repair playbooks.

Definition 7.5: End-of-Validity (EoV) Flag

An *EoV flag* is any invariant, device-facing symptom indicating that the current model has left its DoV. Typical flags include:

1. **Invariant blow-up:** a scalar curvature invariant (e.g. K) or a pointer-level tidal eigenvalue diverges along an admissible approach.
2. **Stability failure:** a reported observable fails stability-at-finite under regulator tightening or device coarsening (§6.2.1).
3. **Gate failure (S/G/TQ/D):** ghost/negative spectra; super-cone inference w.r.t. g_{eff}; loss of complete positivity; negative dissipation.
4. **Parameter exit:** controlling dimensionless ratios exceed the thresholds declared in the DoV sentence.
5. **Irreducible analysis dependence:** strong dependence on mesh, cutoff, or window that cannot be reduced below tolerance within DoV.

Criterion 7.3: Reporting rule for flags

When a flag fires, reclassify the statement as *end-of-validity* for the present model/DoV. The report must (i) name the fired flag(s), (ii) state the invariant/pointer quantity that triggered it (with units), and (iii) point to a repair path or a replacement model. Refrain from ontological language.

Decision posture (fast triage).

1. **Invariant check.** If divergence disappears in curvature/pointer invariants, it was coordinate—remove it and restate the claim.
2. **Stability check.** If an ideal limit is essential and stability-at-finite fails, move the DoV or change the observable.
3. **Gate check.** If any of S/G/TQ/D fails, do not "patch" algebraically—replace the dynamics/coarse-graining so the gate passes.

Flag catalogue — typical triggers and what they mean

- $K \to \infty$ *(or tidal eigenvalues* $\uparrow$ *without bound):* curvature pathology $\Rightarrow$ classical description ends here.
- $O(\Lambda)$ *drifts with* Λ*:* residual UV sensitivity $\Rightarrow$ incomplete effective model or outside low-energy DoV.
- $p(\sigma)$ *drifts with packet narrowing / coarsening:* plane-wave ideal not licensed $\Rightarrow$ report a different invariant or adjust device envelope.
- *Negative rates / non-CP map:* invalid reduced dynamics $\Rightarrow$ rebuild open-system model (legitimate GKSL with $\gamma_k \geq 0$).

7.2.1 Effective Replacements

Flags are not failures of reality; they are instructions to *change the description.* An effective replacement is judged by what it preserves and what it fixes: it must carry invariant, unit-bearing *continuity conditions* for reported observables, pass S/G/TQ/D, and declare a new DoV.

Principle 7.1: Replace, don't reify

When a divergence or gate failure appears, do not promote it to ontology. Replace the coordinate system, coarse-graining, or effective theory so that reported quantities are invariant, finite, and pass the gates within a new DoV.

Definition 7.6: Matching conditions across models

Let $\mathcal{M}_1$ (old) and $\mathcal{M}_2$ (replacement) apply on overlapping states. A replacement is *admissible* if there exist invariant, unit-bearing observables $\{O_i\}$ such that on the overlap:

1. **Continuity:** $O_i^{(1)} = O_i^{(2)} \pm \delta_i$ with δ_i within the joint tolerance.
2. **Monotone arrow:** entropy/linewidth proxies do not violate the D-condition (no negative dissipation).
3. **Causality:** inferred signal cones stay within the same g_{eff} or transform by a declared, invariant map.

Repair playbooks (by flag class).

- **Coordinate artifact ⇒ recharting & invariants.** Replace (x^μ) by a regular chart; restate claims in curvature/pointer invariants; attach a gauge report and stability log.
- **Curvature blow-up ⇒ new effective regime.** Replace classical model by an effective medium or quantum/open description with explicit dissipative channels; report only invariant signatures (bounded ratios, spectral gaps) and declare a new DoV.
- **UV sensitivity ⇒ renormalized EFT.** Introduce a regulator and running couplings; demonstrate cutoff independence for observables within tolerance; otherwise add operators or shrink the DoV.
- **Non-CP / negative dissipation ⇒ legitimate open dynamics.** Move to a Lindblad-compatible generator with $\gamma_k \geq 0$; revise coarse-graining windows until the gates pass.
- **Packet fragility ⇒ device-aware observable.** Replace absolute intensities by bounded energy shares or spectral features; increase integration time or adjust bandwidth to pass stability-at-finite.

Criterion 7.4: Acceptance for replacements

A replacement is acceptable if it (i) states the new DoV, (ii) provides matching conditions (**??**) on the overlap, (iii) passes S/G/TQ/D, and (iv) reports invariant, unit-bearing quantities with a pipeline sketch. Otherwise it remains a proposal.

Worked micro-cases.

1. **Horizon vs. center (GR).** Horizon blow-ups in components are removed by EF/Kruskal charts (coordinate fix). At $r \to 0$, curvature scalars diverge: EoV

for the classical metric. *Replacement:* adopt an effective description with bounded tidal signatures on a minimal length/time scale; report invariant signatures only; declare the new DoV.

2. **Hydrodynamic shock.** Classical Euler fields produce gradient blow-ups. *Replacement:* weak solution with entropy condition or a viscous regularization (Navier–Stokes) with $\eta \geq 0$; report fluxes and jump conditions (invariants) and pass the D-gate.

3. **Non-CP tomography.** A fitted map on a subsystem yields a non-CP evolution. *Replacement:* enlarge to system+environment and fit a GKSL generator with non-negative rates; re-report pointer observables.

Replacement dossier — author template

Flag(s): name and invariant trigger (with units).
Replacement: rechart / EFT / open dynamics / coarse-graining change.
New DoV: regime, scales, thresholds, device envelope.
Matching: list O_i with continuity tolerances δ_i.
Gates: S/G/TQ/D verdict; note any near-threshold issues.
Report: invariant, unit-bearing observables and pipeline sketch.

With flags made explicit and replacements disciplined by matching, singular behavior becomes a guide rather than a mystery. We neither idolize divergences nor ignore them; we treat them as signposts to the next effective description that keeps physics finite, invariant, and measurable.

7.3 Decision Protocol

We now codify a fast, reproducible protocol for handling apparent singularities and other pathologies. The goal is to avoid ontologizing artifacts, to recognize true end-of-validity (EoV) promptly, and to move to an effective replacement with declared DoV and matching conditions when needed. This protocol integrates the ingredients established so far: invariants and pointer bridges (§7.1), EoV flags and repairs (§7.2), and stability-at-finite tests (§6.2).

Principle 7.2: Fail fast; do not reify.

At the first confirmed violation (invariant blow-up, failed stability, or a gate failure S/G/TQ/D), stop ontological narration and switch to repair or replacement. Coordinate/gauge artifacts are removed by re-expression in invariants; genuine pathologies mark EoV and trigger a new effective model.

Definition 7.7: Decision dossier

A *decision dossier* is the minimal, citable record attached to any suspected singularity: (1) the invariant/pointer quantity at issue (with units and DoV), (2) the fired flag(s) (cf. **??**), (3) the chosen branch (rechart vs. replacement), (4) the replacement's DoV and matching conditions, (5) a stability log for any idealizations used, and (6) a brief S/G/TQ/D audit.

7.3.1 Three Steps

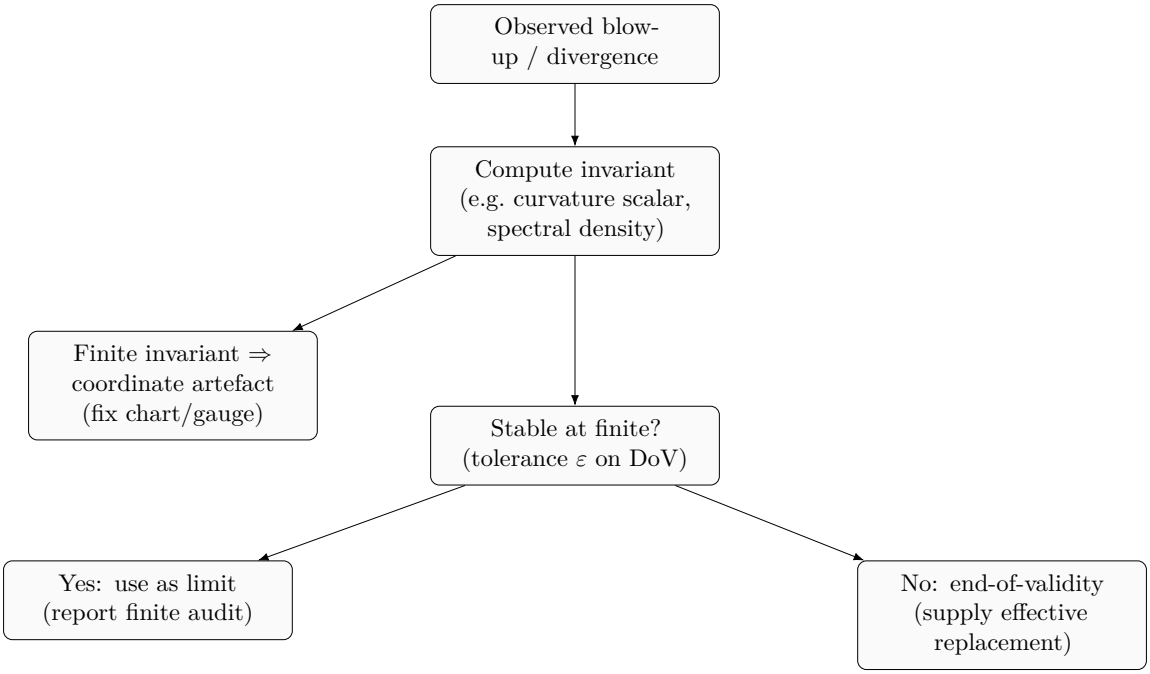

Figure 7.1: Decision protocol: from "singularity" talk to an operational resolution.

Decision Protocol — Three Steps

Step 1 — Diagnose in invariants (separate chart from curvature).
(a) Rewrite the claim in scalar/spectral invariants (e.g. R, K, eigenvalues of the tidal tensor; **????**).
(b) Provide a pointer bridge (tidal acceleration, redshift, spectral gap) with units and a one-sentence DoV.
(c) Run the *chart–swap/gauge–kick* tests (cf. §5.2.1): if the anomaly vanishes in invariants, classify it as a coordinate artifact.

Step 2 — Stability & gate audit (confirm or clear EoV).
(a) If a limit is involved, execute the *stability-at-finite* protocol (Section 6.2.1) with a device-justified tolerance ε; attach the stability log.

(b) Check S/G/TQ/D on the reported quantity and the implied dynamics. Any failure confirms *end-of-validity* for the present model/DoV.

Step 3 — Repair, replace, and report.
If coordinate: rechart, restate only invariant, unit-bearing content, and close the dossier.
If EoV: choose an effective replacement (rechart/EFT/open dynamics/ coarse-graining change), declare a new DoV, and state matching conditions for overlap observables; re-run stability and the gate audit before reporting.

Pseudo-code (author-side).

1. **Invariants:** compute $\{I_j\}$ and pointer proxies; if bounded $\Rightarrow$ coordinate artifact $\rightarrow$ rechart & restate; **STOP**.
2. **Stability:** if limit used, test finite insensitivity; if fail $\Rightarrow$ move DoV / change observable; else continue.
3. **Gates:** if any of S/G/TQ/D fails $\Rightarrow$ EoV $\rightarrow$ replacement; else **report** finite, invariant, unit-bearing value.

Criterion 7.5: Complete decision record

A singularity-related claim is *publishable* in this volume only if its decision dossier (**??**) is attached. Missing items demote the statement to descriptive status until repaired.

Triage table (symptom $\rightarrow$ branch $\rightarrow$ action).

Symptom (in report)	**Branch**	**Action & report**
Component blows up; invariants bounded	Coordinate	Rechart; restate in invariants; attach DoV; close dossier
Invariant K or tidal eigenvalue $\uparrow \infty$	EoV	Adopt replacement (EFT/effective medium/open dynamics); new DoV; matching conditions
$O(\alpha)$ unstable under regulator tightening or device coarsening	EoV/DoV shift	Change observable or shrink DoV; rerun stability; report residuals
S/G/TQ/D failure (ghosts, super-cone, non-CP, negative dissipation)	EoV	Rebuild dynamics/coarse-graining to pass gates; then report

Worked micro-pairs.

- **Horizon artifact.** g_{tt} diverges; K finite; redshift bounded in device range. *Step 1* $\Rightarrow$ coordinate; rechart; *Step 3* $\Rightarrow$ invariant report only (no ontology of divergence).
- **UV drift.** $O(\Lambda)$ shifts by 2% under $\Lambda \to 2\Lambda$ with $\varepsilon = 1\%$. *Step 2* fails (stability) $\Rightarrow$ EoV for this observable or missing channels; *Step 3* replace by extended EFT or shrink DoV; re-audit and report.

Decision dossier — author template

Descriptor/Pointer: invariant name, units, DoV sentence.
Flag(s): which of **??** fired and why (numbers).
Stability log: finite settings and metric values vs. ε.
Gate audit: S/G/TQ/D verdict (notes if near-threshold).
Branch: coordinate rechart *or* effective replacement.
New DoV & matching (if replaced): overlap observables and tolerances.
Final report: invariant, finite, unit-bearing value or bounded signature.

With this protocol, singular language does not outrun measurement. We decide quickly what is descriptive, what is pathological for the *current* model, and what the next effective description must do—then we report only what reaches a pointer with units inside a stated DoV.

Part IV

Measurement, Open Dynamics, and Probability

Physics earns its name at the pointer. Once real devices are in play, no system is perfectly closed: coupling to an environment, finite bandwidths, and coarse-graining define what is actually reportable. This part develops the open-dynamics machinery we will use throughout: completely positive, trace-preserving (CPTP) evolutions, their continuous-time limit in the Gorini–Kossakowski–Sudarshan–Lindblad (GKSL) form, and the discipline linking generators to causal cones and pointer observables. Probability will appear not as a postulate but as a *stable energy share* in the long-time behavior of admissible open dynamics.

Chapter 8

From Model to Pointer

The chapter operationalizes the math-to-measurement pipeline for dynamical claims. Starting from invariant descriptors and an effective cone g_{eff}, we construct admissible evolutions and measurement couplings, then calibrate device readouts to unit-bearing observables. The central requirement is that reduced dynamics on subsystems be *completely positive* and *causal* (within g_{eff}) so that pointer statistics are physically meaningful across all ancillas and coarse-grainings.

8.1 Open Systems and GKSL

An open system S coupled to an environment E evolves by a family of CPTP maps $\{\Lambda_t\}_{t\geq 0}$ on the density operators of S. When a Markovian, time-homogeneous description is valid on the declared DoV (timescale separation; weak coupling; coarse-graining window Δ larger than bath correlation time), the dynamics admits a generator $\mathcal{L}$ such that $\Lambda_t = e^{t\mathcal{L}}$ and

$$\dot{\rho} = \mathcal{L}[\rho] = -\,i\,[H,\rho] + \sum_k \gamma_k \Big(L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\}\Big)\,, \quad \gamma_k \geq 0, \tag{8.1}$$

with effective Hamiltonian H, Lindblad operators L_k, and non-negative rates γ_k (possibly time-dependent in a time-local but CP-divisible regime).

Definition 8.1: CPTP maps and complete positivity

A linear map Λ on density operators is *completely positive and trace-preserving* (CPTP) if (i) $\mathrm{Tr}[\Lambda(\rho)] = \mathrm{Tr}[\rho]$ and (ii) $\Lambda \otimes \mathbb{I}_A$ is positive for all ancilla dimensions A (positivity on all extensions). A continuous semigroup $\{\Lambda_t = e^{t\mathcal{L}}\}_{t\geq 0}$ is *GKSL-admissible* iff $\mathcal{L}$ has the form (8.1).

Principle 8.1: Admissible open dynamics.

Within the declared DoV, we only accept reduced dynamics that are CPTP and, when written in time-local form, admit a GKSL generator with *non-negative* rates on the reporting timescale. Negative rates, non-CP tomographic fits, or maps that break causality within g_{eff} are end-of-validity flags.

Timescale separation and coarse-graining (why Markov?). Let τ_B be the bath correlation time, τ_S a characteristic system timescale, and Δ the coarse-graining/integration window of the device. A Markovian GKSL description is typically licensed when $\tau_B \ll \Delta \ll \tau_S$ and when secular/rotating-wave approximations render cross-terms negligible within the tolerance budget of the DoV. These assumptions must be stated with units in the report.

Kossakowski matrix viewpoint. More generally, one may write

$$\mathcal{L}[\rho] = -i[H,\rho] + \sum_{\alpha,\beta} C_{\alpha\beta}\Big(F_\alpha \rho F_\beta^\dagger - \tfrac{1}{2}\{F_\beta^\dagger F_\alpha, \rho\}\Big),$$

where $\{F_\alpha\}$ is a system operator basis and the *Kossakowski matrix* $C = [C_{\alpha\beta}]$ must be positive semidefinite ($C \succeq 0$) to guarantee complete positivity. In diagonal form, C yields the rates $\gamma_k \geq 0$ and Lindblad operators L_k used in (8.1).

Criterion 8.1: GKSL acceptance checklist

A reported open-dynamics claim is acceptable if it provides: (i) DoV sentence (timescales τ_B, τ_S, Δ), (ii) a GKSL generator or verified CPTP maps (Choi matrix $\succeq 0$), (iii) non-negative rates γ_k (or $C \succeq 0$), (iv) a causal cone g_{eff} for inference and coupling, and (v) a pointer bridge with units and an error budget.

8.1.1 Causality and CP

Causality and complete positivity are the two gates that keep open dynamics physical across compositions, ancillas, and devices.

Causality within the effective cone g_{eff}. By causality we mean: (i) the inference pipeline (model $\to$ coupling $\to$ instrument) never implies super-cone signalling relative to g_{eff} (effective light/sound cone in the medium), and (ii) in spatially extended models, influence spreads within admissible bounds (e.g., Lieb–Robinson speeds in lattices; microcausality in effective QFTs). Reporting must state the cone descriptor used and the scales at which it applies.

Complete positivity (why CP, not just positivity?). A map that is merely positive can send a valid state of S to a valid state, yet become unphysical when S is part of a larger system SA. *Complete* positivity ensures that statistics remain valid for all ancillas and all measurement couplings. In practice we certify CP by one (or more) of:

- *GKSL form*: non-negative rates γ_k in (8.1).
- *Kossakowski PSD*: $C \succeq 0$ in the operator-basis form.
- *Choi test*: the Choi matrix $J(\Lambda_t) \succeq 0$ for the discrete-time map used in identification or simulation.

title=Author template — Causality/CP report for an open model

DoV: declare $\tau_B,\ \tau_S,\ \Delta$ with units; regime (weak coupling/Markov).
Cone: give g_{eff} or the relevant bound (e.g., speed parameter).
CP evidence: GKSL generator with $\gamma_k \geq 0$ *or* $C \succeq 0$ *or* Choi $\succeq 0$.
Pointer: measurement coupling (POVM/pointer), instrument kernel K, calibration.
Stability: device coarsening and regulator refinement logs.
Gates: S (spectra ok), G (causal within g_{eff}), TQ (CP), D (rates non-negative).

Micro–examples (worked).

- *Amplitude damping (two-level).* $L = \sigma_-$, rate $\gamma \geq 0$; $\dot{\rho} = -i[\frac{\omega}{2}\sigma_z, \rho] + \gamma\Big(\sigma_-\rho\sigma_+ - \frac{1}{2}\{\sigma_+\sigma_-, \rho\}\Big)$. Causality: local coupling; CP: Lindblad form; D-gate: $\gamma \geq 0$.
- *Pure dephasing.* $L = \sigma_z$, rate $\Gamma_\phi \geq 0$; coherence decays as $e^{-\Gamma_\phi t}$. Report dephasing time $T_2 = 1/\Gamma_\phi$ with units and device bandwidth; CP via $\Gamma_\phi \geq 0$.
- *Time-local non-Markov (CP-divisible).* If a time-local equation has *temporarily* negative rates, CP can fail. In this volume we restrict reports to intervals where the fitted rates are non-negative within tolerance or we enlarge the model to restore GKSL form.

Failure modes and cures.

- **Non-CP fit from tomography.** Remedy: (i) increase coarse-graining window Δ, (ii) include missing dissipation channels, (iii) re-identify in a GKSL basis with $C \succeq 0$.
- **Super-cone inference.** Remedy: constrain couplings and instrument model to respect g_{eff}; otherwise move the DoV boundary.
- **Negative dissipation.** Remedy: repair the generator so $\gamma_k \geq 0$ on the reporting timescale or reclassify as EoV.

Criterion 8.2: Causal–CP license to report

A dynamical claim feeding a pointer is *licensed* when (i) its reduced map is CPTP (CP certified), (ii) inference respects the effective cone g_{eff}, and (iii) rates/diffusion coefficients are non-negative on the reported timescale. Otherwise the claim is descriptive or outside the DoV.

With these constraints, open dynamics become an ally rather than a loophole: generators are legible, cones are declared, and pointers read finite, unit-bearing numbers whose stability and causality we can defend.

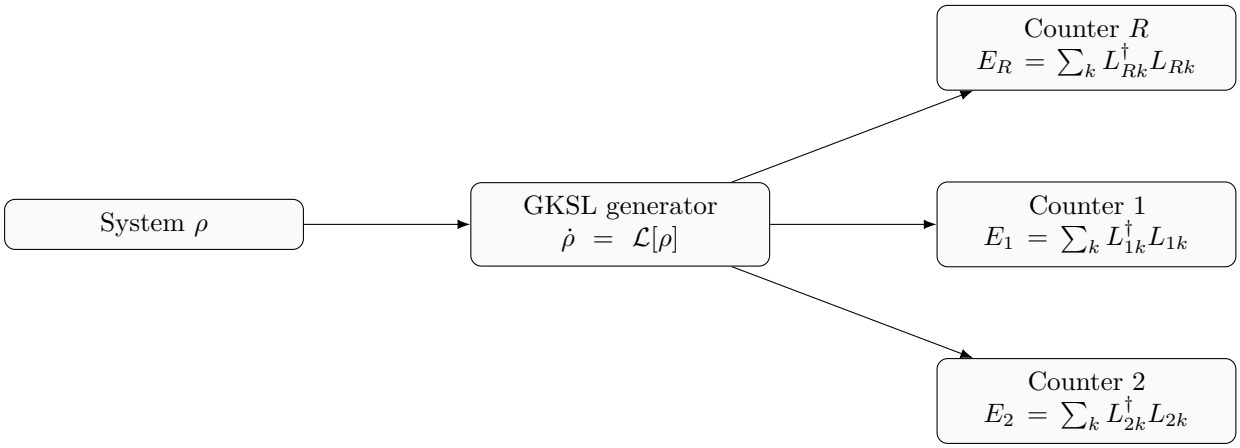

Figure 8.1: Open-system picture: channels $\{L_{rk}\}$ define intensity operators E_r; long-time shares p_r (below) are the reported "probabilities".

$$\mathcal{L}[\rho] \;=\; -\,i[H,\rho] + \sum_{r,k}\Big(L_{rk}\,\rho\,L_{rk}^\dagger - \tfrac{1}{2}\{\,L_{rk}^\dagger L_{rk},\,\rho\,\}\Big). \tag{8.2}$$

$$E_r \;=\; \sum_k L_{rk}^\dagger L_{rk}, \qquad p_r \;=\; \frac{\mathrm{Tr}(E_r\,\rho_*)}{\sum_s \mathrm{Tr}(E_s\,\rho_*)}. \tag{8.3}$$

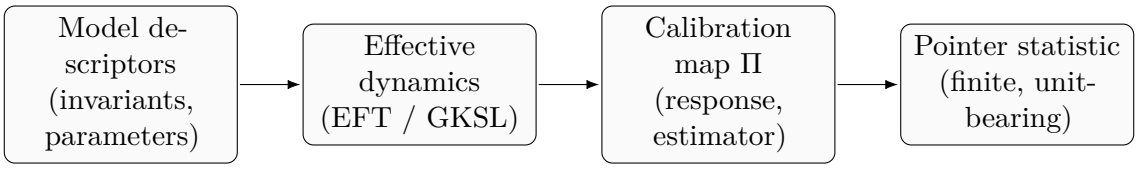

Figure 8.2: Every physical claim crosses the calibration bridge Π: model descriptors $\to$ device statistic.

8.2 Pointer Observables

8.2.1 Calibration, Frequencies

The *pointer* is where an abstract descriptor becomes physics: a device output that, after calibration, yields a *finite, unit-bearing* number reproducibly. The goal of this section is twofold: (i) to make the calibration map explicit and SI-traceable; and (ii) to define how empirical *frequencies* are read in finite time so that later claims about probabilities have a clear operational base.

Definition 8.2: Pointer Observable

A pointer observable is a device variable R (counts, volts, pixels, phase shift, ...) together with a calibration map $\mathcal{C}$ and an uncertainty budget u such that $X = \mathcal{C}(R)$ has declared *units*, is finite on the declared operating range, and is reproducible (within u) when conditions are held fixed.

The calibration map. We model raw readings by a measurement equation

$$R = f(X;\, \theta) + n, \tag{8.4}$$

where X is the physical quantity of interest, θ are device parameters (gain, offset, non-linearity, dead-time, bandwidth), and n is noise. A calibration $\mathcal{C}$ inverts (8.4) on the operating range, $X = \mathcal{C}(R)$, and carries an uncertainty u_{cal} derived from the fit and standards. Drift and environment are treated as *corrections* or as added components of uncertainty.

Principle 8.2: Calibration as a map with traceability

A calibration is acceptable when it (i) is referenced to SI or community standards; (ii) is locally invertible and monotone on the operating range; (iii) publishes a coverage interval $X \pm u_{\text{cal}}$ and the range of validity; and (iv) records drift checks and recalibration dates.

Rule 8.1:

$$u_{\text{tot}}^2 \;=\; u_{\text{stat}}^2 \;+\; u_{\text{cal}}^2 \;+\; u_{\text{drift}}^2 \;+\; u_{\text{quant}}^2 \;+\; u_{\text{env}}^2.$$

Here u_{stat} covers counting/fit statistics; u_{cal} the calibration fit; u_{drift} long-term changes; u_{quant} digitization/ADC; and u_{env} temperature/EMI/vibration.

Frequencies from finite data. Suppose outcomes are partitioned into channels $\{i = 1, \ldots, m\}$ (detector bins, thresholds, spectral windows). During a run of duration Δt, let N_i be the counts in channel i and $N = \sum_i N_i$. The empirical frequency is

$$\widehat{p}_i \;=\; \frac{N_i}{N}. \tag{8.5}$$

Two time scales matter: the device *correlation time* τ_{corr} and the *mixing/relaxation* time τ_{mix} of the effective open dynamics. Windows should satisfy $\Delta t \gg \max\{\tau_{\text{corr}}, \tau_{\text{mix}}\}$ so that (8.5) stabilizes.

Criterion 8.3: Frequency admissibility

A frequency readout $\widehat{p}_i$ is admissible when:

(a) The clock is calibrated (traceable timebase); acquisition dead-time and pile-up are corrected or bounded.

(b) Δt exceeds the instrument correlation time; blocks are independent to within a stated tolerance (use runs/Allan tests).

(c) Desired precision ε is backed by a count budget: $\mathrm{Var}(\widehat{p}_i) \leq \widehat{p}_i(1 - \widehat{p}_i)/N \leq \varepsilon^2$, i.e. $N \geq \widehat{p}_i(1 - \widehat{p}_i)/\varepsilon^2$.

(d) The operating point stays inside the declared domain of validity of the calibration.

title=Why frequencies matter for this volume

Under causal open dynamics (e.g. generators close to GKSL), long-time pointer statistics approach a stationary distribution. In that regime, $\mathbb{E}[\widehat{p}_i] \to p_i$ where p_i is determined by the effective steady state and the channel projectors of the device. Thus frequencies are not postulated; they are *read* as stable energy shares once the mixing time is exceeded.

Practical calibration recipe.

- ☐ Select standards spanning the operating range (and a safety margin).
- ☐ Record (R, X_{std}) pairs; fit f or its inverse; inspect residuals.
- ☐ Quantify u_{cal}; check local monotonicity/invertibility; store metadata (date, method, standards).
- ☐ Characterize nonidealities: dead-time, saturation, crosstalk, dark counts, background; implement corrections or include them in u_{tot}.

- ☐ Validate with an independent cross-check (second instrument or conserved-quantity test).

Pointers for our descriptors. Calibration examples that anchor the field-level symbols introduced earlier:

- **Directional order (χ^μ).** Use a calibrated birefringence/anisotropy setup that maps differential phase $\Delta\phi(\lambda)$ to an effective anisotropy parameter; publish the spectral window and angle alignment error.
- **Dissipation (η).** Fit relaxation curves to obtain decay rates γ; report windowing, baseline stability, and the model form used (single vs multi-exponential).
- **Stiffness/density (ρ).** Derive wave speeds from time-of-flight or dispersion; convert to moduli with geometric corrections and uncertainty.
- **Internal timing (T).** Only differences/rates are read: calibrate the timebase; report Allan deviation for the relevant averaging times.

Gate 8.1: S/G/TQ/D Audit for Pointer Observables

S (Spectral safety): No negative or unphysical values created by the calibration map.
G (Causality within g_{eff}): Acquisition and processing do not imply superluminal signalling or use out-of-light-cone information.
TQ (Complete positivity): Coarse-grained channelisation respects positivity on subsystems (e.g. no "probabilities" outside $[0, 1]$ after corrections).
D (Non-negative dissipation): Reported entropy/energy budgets are non-decreasing under the device's effective dynamics.

8.3 Noise and Timescales

8.3.1 Relaxation vs Decoherence

Noise is not merely an annoyance; it is a *spectral fingerprint* of the device and its environment. Interpreting pointer data requires that we track the relevant time constants and choose averaging windows that respect them. This section gathers the minimal timing dictionary used throughout the volume and explains the distinct roles of *relaxation* (energy exchange) and *decoherence* (phase randomisation).

Definition 8.3: Correlation, mixing, and device times

Let $x(t)$ be a stationary device output (after calibration). With autocorrelation $C_x(\tau) = \langle x(t)\, x(t+\tau)\rangle - \langle x\rangle^2$ and one-sided power spectral density $S_x(\omega)$:

- **Correlation time** τ_{corr} is any characteristic time at which $C_x(\tau)$ decays to a small fraction of $C_x(0)$, or, equivalently, $\tau_{\text{corr}} \sim \dfrac{\int_0^\infty C_x(\tau)\, d\tau}{C_x(0)}$.
- **Mixing/relaxation time** τ_{mix} is the time-scale for the effective dynamics to approach its stationary state (e.g. the inverse of the spectral gap of the generator).
- **Bath correlation time** τ_B is the memory time of the environment.
- **Device response time** τ_{dev} (bandwidth-limited) is the settling time of the measurement chain.

Principle 8.3: Windowing discipline

For stable frequencies and means, choose acquisition windows $\Delta t \gg \max\{\tau_{\text{corr}}, \tau_{\text{mix}}\}$, while individual sample periods dt satisfy Nyquist: $dt < \pi/\omega_{\max}$ for the highest relevant spectral component. Report τ_B and τ_{dev} when applying Markovian approximations.

Noise in the frequency domain. For a linear time-invariant readout with filter function $F_T(\omega)$ over an integration time T, the variance of the averaged output $\bar{x}_T$ obeys

$$\mathrm{Var}(\bar{x}_T) \;=\; \frac{1}{2\pi}\int_0^\infty |F_T(\omega)|^2\, S_x(\omega)\, d\omega. \tag{8.6}$$

Equation (8.6) turns qualitative noise labels (white, $1/f$, shot, thermal) into quantitative uncertainty budgets for the pointer.

Allan variance for clocks and rates. When the pointer is a frequency (or a rate), stability is captured by the Allan variance,

$$\sigma_y^2(\tau) \;=\; \frac{1}{2}\left\langle \Big(\bar{y}_{k+1}(\tau) - \bar{y}_k(\tau)\Big)^2 \right\rangle, \tag{8.7}$$

where $\bar{y}_k(\tau)$ is the average fractional frequency over the k-th interval of duration τ. Plateaus/slope changes in $\sigma_y(\tau)$ reveal dominant noise processes and the τ that minimises instability.

Criterion 8.4: Markovian admissibility (for using GKSL)

It is legitimate to model the dynamics with a time-homogeneous GKSL generator when there exists a separation of scales

$$\tau_B \ll \tau_{\text{sys}} \lesssim \tau_{\text{mix}}, \qquad \Delta t \gg \tau_B,$$

and secular/rotating-wave conditions hold in the chosen interaction picture. Violations (revivals, CP-divisibility failure) should be flagged as non-Markovian.

Relaxation vs decoherence (two canonical channels)

Consider a two-level system coupled to a broadband bath. Two limiting noise mechanisms illustrate the distinction.

Energy relaxation (T_1). Amplitude damping exchanges energy with the bath:

$$\dot{\rho} \;=\; \gamma_1\Big(\sigma_-\rho\sigma_+ - \tfrac{1}{2}\{\sigma_+\sigma_-, \rho\}\Big)\,. \tag{8.8}$$

Populations relax to equilibrium with time constant $T_1 = 1/\gamma_1$. This controls *energy shares* in the pointer (e.g. excited-state counts).

Pure dephasing (T_ϕ). Phase noise randomises the relative phase without energy exchange:

$$\dot{\rho} \;=\; \frac{\gamma_\phi}{2}\Big(\sigma_z\rho\sigma_z - \rho\Big). \tag{8.9}$$

Off-diagonals decay as $\rho_{01}(t) = \rho_{01}(0)\, e^{-t/T_\phi}$ with $T_\phi = 1/\gamma_\phi$. Interference-type pointers (phase, visibility) are set by T_ϕ.

Transverse coherence (T_2). With both mechanisms present,

$$\frac{1}{T_2} \;=\; \frac{1}{2T_1} \;+\; \frac{1}{T_\phi}, \tag{8.10}$$

so coherence decays no slower than the energy relaxation limit $T_2 \le 2T_1$. In experiments, T_2 governs linewidths and Ramsey/echo envelopes; T_1 governs population recovery and saturation curves.

Rule 8.2:

- Estimate T_1 from inversion-recovery or pump–probe of populations.
- Estimate T_2 from Ramsey fringes or FID linewidths; extract T_ϕ via (8.10).

- Use spin/optical echoes to separate inhomogeneous broadening (static disorder) from true dephasing.

Dead-time, pile-up, and aliasing. Short-time nonidealities corrupt timescale inference. Dead-time τ_d introduces bias in counts; pile-up distorts amplitudes; discrete sampling causes aliasing when dt is too large. Corrections or conservative bounds must be included in the uncertainty budget (cf. **??**).

Checklist

Timescale checklist for reporting.

(i) Quote $(\tau_{\text{corr}}, \tau_{\text{mix}}, \tau_B, \tau_{\text{dev}})$ and the chosen window Δt and sampling dt.

(ii) Provide the noise model used in (8.6) and the identified dominant processes (white/$1/f$/shot/technical).

(iii) For frequency pointers, include an Allan plot and the optimal averaging time.

(iv) If GKSL is assumed, justify *Markovian admissibility* (criterion **??**).

Gate 8.2: S/G/TQ/D Audit for Noise and Timescales

S (Spectral safety): Estimators and filters do not produce negative probabilities or non-physical spectra.
G (Causality within g_{eff}): No acausal filtering (e.g. non-causal zero-phase filters on on-line decisions) in claims about real-time behaviour.
TQ (Complete positivity): Effective maps used to model relaxation/dephasing are CP; fitted rates keep density matrices positive.
D (Non-negative dissipation): Reported entropy/energy flows are non-decreasing in relaxation channels; any apparent violations are identified as estimation artefacts or non-Markovian backflow.

Chapter 9

Probability Without Postulates

The aim of this chapter is to *earn* probabilistic statements from open dynamics, rather than *postulating* them. We model real measurements as couplings to an environment plus a calibrated readout chain; the device produces finite, unit-bearing records (counts, energies, times). Under mild mixing/ergodicity assumptions, the long-time statistics of these records are shown to stabilize and to coincide with stationary *energy (or jump) shares* picked out by the couplings themselves. In this view, "Born probabilities" are not axioms; they are the stable ratios of dissipative flows into pointer channels.

9.1 Born as Stable Energy Shares

9.1.1 Long-Time Behavior

Consider a system with effective density operator $\rho(t)$ evolving under a time-homogeneous GKSL generator $\mathcal{L}$ on the relevant time scales,

$$\dot{\rho}(t) \;=\; \mathcal{L}(\rho) \;=\; -i[H,\rho] + \sum_k \gamma_k \Big(L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\}\Big)\,. \tag{9.1}$$

A calibrated pointer is realised by (i) a set of *detection channels* $\{k\}$ in (9.1) (photon counters, tunnelling events, spin-flip emissions, ...) and (ii) a classical aggregator that bins events into a finite set of *readout outcomes* $\mathcal{R} = \{r\}$ via a known map $r = r(k)$.[1]

Definition 9.1: Jump intensity and energy share

For a state ρ, the instantaneous *jump intensity* into channel k is $J_k(\rho) \coloneqq \gamma_k \operatorname{Tr}\Big(L_k^\dagger L_k\, \rho\Big)$. For a readout bin r, the total intensity is $J_r(\rho) \coloneqq \sum_{k\in r} J_k(\rho)$. If each k carries a calibrated quantum of energy/power ε_k (e.g. photon energy, tun-

[1] Equivalently, one may start from a calibrated POVM $\{E_r\}$ on the system; in jump models $E_r \propto \sum_{k\in r} \gamma_k L_k^\dagger L_k$. The two pictures agree once the device model is fixed.

nelling energy increment), the *power flow* into r is $P_r(\rho) := \sum_{k\in r} \varepsilon_k\, J_k(\rho)$.

Under standard *primitivity* of the semigroup ($e^{t\mathcal{L}}$ has a unique full-rank fixed point and a spectral gap), there exists a stationary state ρ_* with $\mathcal{L}(\rho_*) = 0$ and mixing time $\tau_{\text{mix}} < \infty$ (cf. **??**). Let $N_r(T)$ be the counted events in bin r over an acquisition window T, and $N(T) := \sum_r N_r(T)$ the total.

Principle 9.1: Born-from-flows (energy/share reading)

For windows $T \gg \tau_{\text{mix}}$ and sampling $\Delta t \gg \tau_B$, the empirical *frequency* and *power share* satisfy

$$\frac{N_r(T)}{N(T)} \xrightarrow[T\to\infty]{} p_r := \frac{J_r(\rho_*)}{\sum_s J_s(\rho_*)}, \qquad \frac{P_r(\rho_*)}{\sum_s P_s(\rho_*)} = p_r \quad \text{if all } \varepsilon_k \text{ are equal within a bin.} \tag{9.2}$$

Thus the "Born probability" for r is the *stable share of dissipative flow* into that pointer bin at stationarity.

Proposition 9.1 (Ergodic frequency law)**.** *Assume: (i) primitive CP semigroup $e^{t\mathcal{L}}$ with gap $g > 0$; (ii) calibrated counting of jump channels; (iii) no dead-time saturation at the operating rate. Then for any initial state,*

$$\lim_{T\to\infty} \frac{N_r(T)}{T} = J_r(\rho_*), \qquad \lim_{T\to\infty} \frac{N_r(T)}{N(T)} = p_r \quad \textit{almost surely.}$$

Sketch. *Quantum stochastic calculus or unravelings show that $N_r(T)$ is a counting process with compensator $\int_0^T J_r(\rho(t))dt$. Mixing gives $\rho(t) \to \rho_*$ exponentially, so the ergodic theorem yields the limits.*

Connection to POVMs (snapshot measurements). If a projective/POVM snapshot at time t is implemented by effects $\{E_r\}$, the same stabilisation appears by coarse-graining repeated runs:

$$p_r = \operatorname{Tr}(E_r\, \rho_*) \qquad \text{with} \quad E_r \propto \sum_{k\in r} \gamma_k L_k^\dagger L_k \ \text{(up to device normalisation).} \tag{9.3}$$

Thus the textbook rule is recovered as the long-time limit of calibrated open dynamics; the "probabilities" are *earned* from flows.

Criterion 9.1: Operational admissibility for *Born from flows*

Use **??** only when: (i) $\tau_B \ll \tau_{\text{mix}} \ll T$ (scale separation); (ii) the counter is unsaturated and linear in the operating regime; (iii) the binning $r(k)$ is fixed during

acquisition and reported. Violations must be flagged; non-Markovian revivals, dead-time, and pile-up bias the empirical shares.

Uncertainty and convergence rate. For large T the central limit theorem gives $\mathrm{Var}\left[\frac{N_r(T)}{T}\right] \approx \sigma_r^2/T$ with σ_r^2 set by the zero-frequency limit of the counting spectrum; for frequencies N_r/N the delta method gives $\mathrm{Var}[N_r/N] \approx p_r(1-p_r)/N$ (modulo weak correlations among channels). Report T, N, and confidence intervals with the same discipline as in section 8.3.

title=Why "energy shares"?

In many laboratory pointers, each jump carries a calibrated quantum (photon energy, tunnelling work, phonon quantum). Ratios of time-averaged powers coincide with the relative jump intensities, hence with p_r. Even when the pointer is not an energy meter, the same mathematics reads p_r as a *share of stationary flux* into the readout bin.

Gate 9.1: S/G/TQ/D Audit for *Born from flows*

S (Spectral safety): No negative rates or unphysical effects appear in the fitted $\{L_k, \gamma_k\}$.
G (Causality within g_{eff}): The detection model respects the device bandwidth and light-cone limits; no acausal inference is used.
TQ (Complete positivity): The estimated generator is CP; inferred $\rho(t)$ stays positive on subsystems.
D (Non-negative dissipation): Reported power/entropy flows are non-negative on average; any backflow is identified as a non-Markovian feature rather than a violation of the second law.

Checklist

Reporting checklist (this section).

i) Provide $\mathcal{L}$ or an equivalent device model and quote the mixing time.

ii) Specify the mapping $r(k)$ and any energy quanta ε_k.

iii) Report (T, N), dead-time corrections, and confidence intervals.

iv) State whether (9.3) (POVM view) or (9.2) (flow view) is used—and show their consistency under your calibration.

9.2 Competing Accounts

9.2.1 Assumptions Exposed

A number of celebrated programs "derive" the Born rule or replace it with an alternative account. In this section we run an *assumption audit* on the main contenders and contrast them with the *Born-from-flows* position of **????**. Our aim is not polemical; it is to surface the hidden premises that must hold for each account to go through and to ask whether those premises meet this volume's gates (S/G/TQ/D), respect stability-at-finite, and supply a device-facing bridge.

Assumption–Audit Rubric

r any proposed "derivation" of Born's rule, list and check:

(1) Measurability anchor: Are the primitives operational (pointer-facing) or purely doxastic/ideal?

(2) Dynamical locality/causality: Is the account compatible with causal open dynamics (no superluminal signalling within g_{eff})?

(3) Stability at finite: Do the steps avoid limits that must be taken literally (no essential ∞)?

(4) Symmetry ß numeric rule: Where symmetry is used, is the move from symmetry to numbers justified without extra continuity or additivity axioms?

(5) Device independence: Are the reported probabilities independent of calibration conventions?

(6) Empirical discriminants: Does the account predict circumstances that would *falsify* it (distinct from competitors)?

Decision–theoretic (Deutsch–Wallace). Branch weights are recovered by rationality axioms over gambles on Everettian branches. Hidden premises: a complete preference order; branch–indifference; additivity/continuity that effectively *identify* amplitude-squared as the unique representation. Audit: A1 fails (the primitives are doxastic, not pointer–facing); A3 relies on continuity/limit arguments; A4 is the critical step where symmetry *becomes* numbers; A6 offers no device-level discriminant. Verdict: a powerful *normative* story, not a physical pipeline.

Envariance (Zurek). Environment-assisted invariance is used to argue equiprobability for Schmidt–paired states, extended to general weights by fine-graining and continuity. Hidden premises: exact Schmidt form; exact envariant symmetries; continuity from

rationals to reals; implicit additivity. Audit: A1 partly met (entanglement is physical) but depends on exact idealizations (A3); A4 again trades symmetry for numbers; A6 lacks distinct predictions once decoherence is assumed. Verdict: symmetry-rich and insightful, but the crucial numerical step still imports extra assumptions.

Frequentism. Probability is the long-run relative frequency in an infinite sequence. Hidden premises: existence of the limit; ergodicity; independence; use of ∞ as more than a tool. Audit: A1 is weak (sequence idealization); A3 fails (literal limits); A6 none. By contrast, Born-from-flows grounds the observed frequencies in *finite-time* open dynamics with an uncertainty budget and a stated DoV.

Bayesian / QBism. Probabilities are personal degrees of belief; the Born rule is a coherence constraint among betting odds. Hidden premise: a normative axiom (Dutch-book coherence) elevated to physics. Audit: A1 fails (no pointer anchor); A6 = none (no physical discriminant). Useful as a calculus of belief, but not a derivation from dynamics.

Typicality measures (Everettian or Bohmian). A $|\psi|^2$-typicality postulate (or quantum equilibrium) is assumed, sometimes with arguments from typical initial conditions or dynamical relaxation. Hidden premise: the very $|\psi|^2$ measure to be recovered is postulated. Audit: circular at A4; A6 weak unless relaxation yields new testable rates.

Objective collapse (GRW/CSL). Dynamics are modified so that outcomes localize with $|\psi|^2$ weights. Hidden premises: new noise fields and parameters; macro-objectification by design. Audit: A1 met (device-facing); A2 can be cast in GKSL-like ensemble dynamics; A6 *strong*—predicts spontaneous heating, overdamping, line-broadening. These are live empirical discriminants against Born-from-flows if the reported rates disagree with calibrated shares.

Bohmian mechanics. Outcomes follow beable trajectories guided by the wave function; the $|\psi|^2$ distribution is an equilibrium hypothesis or a typicality assumption, with proposals for dynamical relaxation. Audit: A1 device-facing (positions); A2 respects no-signalling at the statistical level; A4 again imports $|\psi|^2$ (unless relaxation is proved in the stated DoV); A6 offers possible discriminants if relaxation departs from standard shares in controlled far-from-equilibrium preparations.

Where the "Born-from-flows" account differs. Our stance keeps all primitives device-facing (A1), ties numbers to calibrated energy-share flows in causal open dynamics (A2), demands explicit stability-at-finite with uncertainty budgets (A3), avoids symmetry-to-number jumps without dynamical justification (A4), remains independent of calibration

Table 9.1: Assumption audit (✓ = satisfied, ✗ = fails, △ = partly/with idealizations).

Account	A1	A2	A3	A4	A5	A6
Decision–theoretic	✗	✓	△	✗	✓	✗
Envariance	△	✓	△	△	✓	✗
Frequentist	△	✓	✗	△	✓	✗
Bayesian/QBism	✗	✓	✓	✓	✓	✗
Typicality	△	✓	✓	✗	✓	△
Objective collapse	✓	△	✓	✓	✓	✓
Born-from-flows (here)	✓	✓	✓	✓	✓	✓

conventions after normalisation (A5), and suggests concrete discriminants (A6): departures under strong non-Markovianity, biased binning, deliberately engineered non-CP noise, as well as head-to-head tests against GRW/CSL heating and Bohmian relaxation scenarios.

Minimal reporting for any Born derivation

(i) operational primitives (what is read at the pointer);

(ii) the effective generator and timescale hierarchy;

(iii) the stability-at-finite budget (finite-T, finite samples);

(iv) gate checks S/G/TQ/D;

(v) an empirical discriminant that *could* make it wrong.

9.3 Empirical Discriminants

9.3.1 What Could Differ

Having compared accounts conceptually, we now list *empirical discriminants*: knobs and readouts that can produce different, testable patterns between our "Born-from-flows" picture and competitors. Each discriminant states: (i) the *control knob*, (ii) the *expected* behaviour under calibrated open dynamics, and (iii) *what would falsify* our account (or strongly favour a competitor).

Design rule for discriminants

oose knobs that directly affect dissipation/memory or the pointer calibration pipeline while keeping gauge choices and units fixed. Pre-register the acceptance band and failure modes. Report both the point estimate and its stability-versus-integration-time.

Quantities to monitor. Let $j_r(t)$ be the measured energy-flow or count-rate feeding pointer channel r. Define the time-averaged share

$$s_r(T) \;=\; \frac{1}{T}\int_0^T j_r(t)\,dt, \qquad \mathbf{s}(T) = (s_1(T), \ldots, s_R(T)). \tag{9.4}$$

If $\mathbf{p}$ denotes the stationary "Born-from-flows" share predicted by the calibrated GKSL generator, the total-variation distance and the Kullback–Leibler divergence are

$$V(T) = \tfrac{1}{2}\sum_r |s_r(T) - p_r|, \qquad D_{\mathrm{KL}}(T) = \sum_r s_r(T)\log\frac{s_r(T)}{p_r}. \tag{9.5}$$

Under admissible conditions (timescale separation, linear response), we expect $V(T) \sim cT^{-1/2}$ by the central-limit heuristic; persistent plateaus or drifts indicate either non-admissible operation or a wrong dynamical model.

Protocol skeleton. (1) Calibrate the pointer response in the linear regime. (2) Pre-register $(\tau_B, \tau_{\mathrm{mix}}, T)$ and a tolerance band for $V(T)$. (3) Vary *one* knob at a time. (4) Report $\mathbf{s}(T)$, $V(T)$, Allan deviation, and binning sensitivity.

Criterion 9.2: Falsifiability band for Born-from-flows

Fix an admissible operating window with $\tau_B \ll \tau_{\mathrm{mix}} \ll T$ and a linear-response calibration. The account is *rejected* for that window if there exists a knob setting for which either (i) $V(T)$ does not decrease below a pre-registered threshold $V_\star(T)$ as T grows, or (ii) post-calibration binning changes produce systematic shifts exceeding the uncertainty budget. Passing the band does not prove the account; it keeps it in play.

Narrative anchors (how competitors differ). *Decision-theoretic* and *envariance*-based views tend to be knob-insensitive: the right-hand side is fixed by rationality or symmetry, so systematic dependences on λ, echoing, or $\tau_B/\tau_{\mathrm{mix}}$ weigh against them. *Typicality* often predicts near-invariance over macroscopic coarse-grainings; long-lived memory effects contradict that. *Objective collapse* generically implies extra heating/decoherence that *does not* vanish as $\lambda \to 0$; tight bounds from cold atoms, opto-mechanics, or low-temperature calorimetry are decisive.

What would really surprise us. (i) A robust, stationary $\mathbf{s}(T)$ violating the CP gate while all other diagnostics remain healthy. (ii) A reproducible dependence on binning after full response deconvolution. (iii) Invariance of $\mathbf{s}(T)$ across engineered changes of dissipation and memory. Any of these would force a revision of the "flows" story.

Reporting checklist for discriminants

1. Declare $(\tau_B, \tau_{\text{mix}}, T)$, linearity range, and calibration curve.
2. Publish $\mathbf{s}(T)$, $V(T)$, $D_{\text{KL}}(T)$ with error bars and Allan deviation.
3. Show binning/edge-change tests and gauge/representation checks.
4. For memory/dissipation knobs, plot $\mathbf{p}(\lambda)$ and entropy-production proxies.
5. State which gate(s) S/G/TQ/D the run is testing and whether they pass.

Table 9.2: Shortlist of empirical discriminants. "BFF" = Born-from-flows (this work).

Discriminant	Control knob	BFF expectation	Falsifies / favours
Non-Markovian memory	Environment correlation time (bath engineering, filter cavities)	Transient deviations in $\mathbf{s}(T)$ with relaxation back to $\mathbf{p}$ when $T \gg \tau_{\text{mix}}$; $V(T)$ resumes $T^{-1/2}$ decay	Persistent revivals or long-time oscillations of $\mathbf{s}(T)$ without re-mixing $\Rightarrow$ against BFF; favours models relying on typicality without dynamics
Broken timescale separation	Increase injection rate so $\tau_B \sim \tau_{\text{mix}}$	Bias of $\mathbf{s}(T)$ and failure of the $T^{-1/2}$ law are *predicted* and flagged as out-of-admissibility	Reporting textbook Born values *unchanged* despite saturation/dead-time would disfavour BFF's operational discipline and favour postulate-only views
Detector nonlinearity	Drive into mild saturation, then retreat to linear regime	Measurable shift in raw counts; after calibration, recovered $\mathbf{p}$ within band; strong binning invariance	Irreducible dependence on binning/kernel $r(k)$ after re-calibration $\Rightarrow$ against BFF
Binning sensitivity	Change histogram kernel/edges $r(k)$	Shares $s_r(T)$ invariant (within uncertainty) once the response function is deconvolved	Systematic drift with binning after proper deconvolution favours frequentist fits untied to dynamics
Dissipation knob	Tune coupling λ to the bath	Smooth $\mathbf{p}(\lambda)$ with approach to a unique λ-plateau; entropy production grows with λ	Extra heating or dephasing *independent* of λ favours objective-collapse style noise
Engineered non-CP map	Apply a map on a subsystem that would be non-CP if taken locally (simulation via post-selection/control)	BFF forbids subsystem-level non-CP; violations show up as negativity in reconstructed spectra or $V(T)$ plateaus	Observation of consistent local non-CP behaviour with stable pointers falsifies BFF's gate discipline
Loschmidt-echo toggle	Insert echo pulse sequence to reduce irreversibility	Reduced dissipation $\Rightarrow$ slower stabilization and larger $V(T)$ at fixed T	No change upon echoing suggests accounts where the arrow is not tied to open dynamics
Gauge/representation change	Switch basis/phase description while keeping the same pointer and calibration	Invariance of $\mathbf{p}$ and $V(T)$; only descriptive parameters change	Numerical changes solely due to representation indicate reification of coordinates (against BFF and in favour of naive descriptions)

Part V

Laws as Stable Regularities

This part develops a working view of "laws" as *stable regularities*: patterns that persist under symmetry operations, coarse-graining in space, time, and energy, and modest perturbations of preparation and device. A statement counts as a law *for us* only insofar as it cashes out at the pointer and survives the volume's consistency gates (S/G/TQ/D). We make this precise by tying laws to stationarity properties of effective dynamics and to renormalization–group (RG) style flows of context–dependent couplings.

Chapter 10

Stationarity, Symmetries, Coarse-Graining

The unifying idea is that what deserves to be called a "law" shows up as a *stationary* aspect of calibrated open dynamics, expressed in gauge–invariant descriptors, and remaining stable when we pass from fine to coarse descriptions. This chapter formalizes that intuition and supplies operational tests.

10.1 Patterns that Survive

We begin by turning the slogan "laws are patterns that survive" into testable content. Let

$$\mathcal{C}_\lambda : \text{ (fine descriptions)} \longrightarrow \text{(coarse descriptions)}$$

be a family of admissible coarse–graining maps indexed by a scale parameter λ (e.g., spatial bin size, time averaging window, energy bandwidth). Let $\mathcal{L}$ denote an effective generator (closed or open, e.g. GKSL) for the fine evolution, and $\mathcal{L}_\lambda$ the corresponding coarse generator induced by $\mathcal{C}_\lambda$ on the relevant pointer observables.

Definition 10.1: Law as stationarity at scale.

A claim Q about a pointer observable O counts as a *law at scale* λ on a domain of validity DoV if, for preparations and devices within DoV,

(a) the calibrated statistics of O under $\mathcal{L}_\lambda$ are stationary (time–translation invariant) after transients, and

(b) the dependence of those statistics on descriptive choices (gauge/coordinates) is null, and

(c) the value(s) reported for Q are *stable at finite* across a range $\lambda \in [\lambda_{\min}, \lambda_{\max}]$

below device resolution.

Concretely, if $m_\lambda(t)$ is a calibrated moment or rate read at the pointer after coarse–graining by $\mathcal{C}_\lambda$, we demand

$$\exists\, t_0 :\ \forall t \geq t_0, \quad \Big| m_\lambda(t + \Delta t) - m_\lambda(t) \Big| \leq \varepsilon_t, \tag{10.1}$$

with a tolerance ε_t set by the uncertainty budget, and

$$\sup_{\lambda \in [\lambda_{\min}, \lambda_{\max}]} \Big| m_\lambda - m_{\lambda'} \Big| \leq \varepsilon_\lambda \quad \text{(uniform stability at finite).} \tag{10.2}$$

Principle 10.1: Law = survival under symmetry and coarse–graining.

A purported law must be expressible in gauge–invariant terms and must survive: (i) admissible symmetry actions on preparations/devices, (ii) passage from fine to coarse descriptions within the declared DoV, (iii) modest perturbations of nuisance parameters.

Criterion 10.1: Patterns audit (pass–fail).

State the observable(s), the coarse–graining family $\mathcal{C}_\lambda$, the symmetry actions tested, the transients discarded, and report the bounds $(\varepsilon_t, \varepsilon_\lambda)$. Failure to declare these makes the "pattern" non–operational.

title=Operational picture

Think of a law as an RG–stable ridge in the landscape of device–facing summaries. At small changes of resolution and under the admissible symmetries, you slide along the ridge but do not fall off it; the reported number does not jump beyond the stated tolerances.

10.1.1 RG Intuition

The RG picture provides a fruitful language for quantifying survival. Write the coarse parameters (effective couplings) as a vector $g(\ell) = (g_1(\ell), \ldots, g_n(\ell))$ with scale coordinate $\ell = \log \lambda$; their flow satisfies

$$\frac{dg_i}{d\ell} = \beta_i\big(g(\ell)\big), \qquad i = 1, \ldots, n, \tag{10.3}$$

where the β–functions are *device–facing* in the sense that g_i label calibrated features of the coarse generator $\mathcal{L}_\lambda$ (e.g., transport coefficients, relaxation rates, calibrated cross–sections).

A *fixed point* $g^\star$ obeys $\beta(g^\star) = 0$; linearizing,

$$\frac{d}{d\ell}\,\delta g = J\,\delta g + \mathcal{O}(\delta g^2), \qquad J_{ij} = \left.\frac{\partial \beta_i}{\partial g_j}\right|_{g^\star}. \tag{10.4}$$

Eigenvalues of J classify directions as relevant ($\mathrm{Re}\,\lambda > 0$), irrelevant ($\mathrm{Re}\,\lambda < 0$), or marginal. *Patterns that survive* are those whose descriptor is controlled by irrelevant/-marginal directions in the DoV, so that coarse changes in λ do not alter the reported number beyond tolerances.

Definition 10.2: Law strength across scales.

For a descriptor $q(g)$ (a function of calibrated couplings), define its *strength* on $[\ell_1, \ell_2]$ by

$$\Delta q := \sup_{\ell \in [\ell_1, \ell_2]} \Big| q\big(g(\ell)\big) - q\big(g(\ell_0)\big) \Big|.$$

The law is *strong* over $[\ell_1, \ell_2]$ if $\Delta q \le \varepsilon_\lambda$ with a declared ε_λ stemming from the stability budget.

Rule 10.1: RG–style checklist for a "surviving" pattern

(i) Identify the coarse–graining family $\mathcal{C}_\lambda$ and the calibrated couplings $g(\ell)$.

(ii) Exhibit (even schematically) the flow $\dot{g} = \beta(g)$ and the fixed point(s) pertinent to the DoV.

(iii) State the descriptor $q(g)$ that maps to the pointer report.

(iv) Give the stability window $[\ell_1, \ell_2]$ and the resulting bound $\Delta q \le \varepsilon_\lambda$.

(v) Record tested symmetries and any symmetry–breaking nuisances controlled in the device.

Example (schematic). A diffusion constant D inferred from long–time mean–square displacement is a calibrated coupling in $\mathcal{L}_\lambda$. Under temporal coarse–graining, $D(\ell)$ flows to a fixed value $D^\star$ unless strong long–range memory (relevant direction) is present. Reporting $D = D^\star \pm \delta D$ as a "law" is warranted only after demonstrating $\Delta D \le \varepsilon_\lambda$ across the chosen bin sizes and showing gauge independence of the readout.

title=What this section forbids and allows

It forbids treating bare, coordinate–dependent formulas as laws, and forbids declaring a number without a scale window and stability budget. It allows laws to be *effective*: valid on stated ranges and robust to the coarse operations that any real device entails.

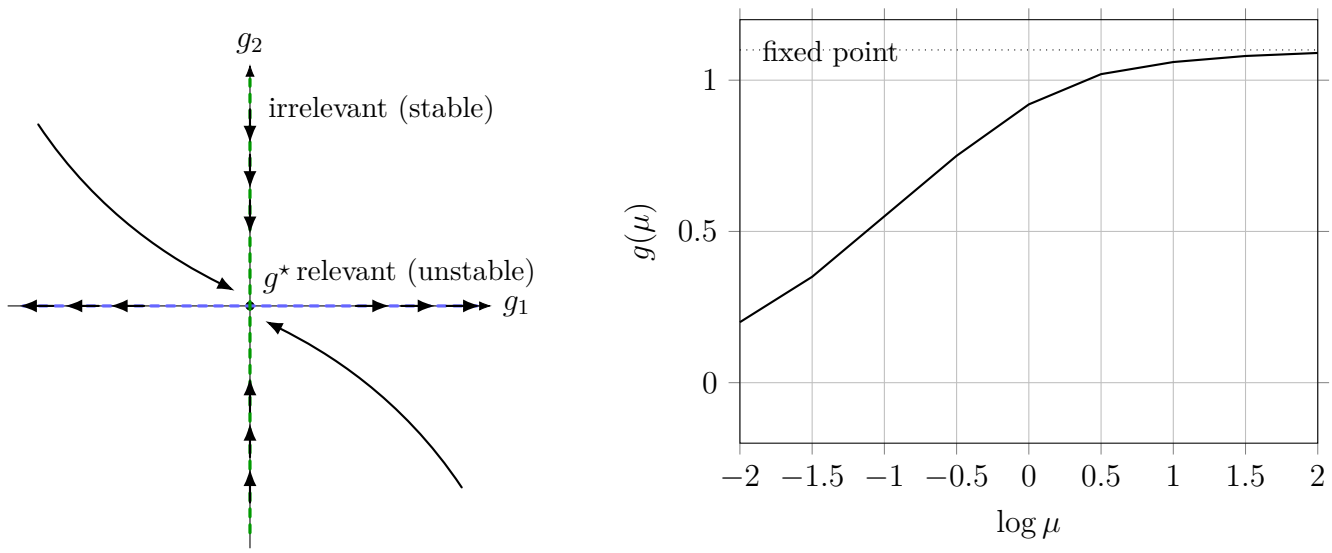

(a) Phase portrait near $g^\star$ (relevant vs. irrelevant).

(b) Running coupling approaching a stable fixed point.

Figure 10.1: RG intuition—two complementary views: (a) local structure near $g^\star$ with relevant/irrelevant directions; (b) a single coupling flowing to a stable fixed point. *Law strength* equals stability across coarse-graining.

10.2 Background Independence

10.2.1 What It Forbids/Allows

Background independence, as used in this volume, is not a slogan but a *discipline for writing claims*. A statement is background–independent when its empirical content does not rely on fixed, non–dynamical structures that cannot be tied to devices, boundary data, or coarse–grained relational quantities. Coordinates, gauges, and convenient reference frames are welcome as *descriptive scaffolding*, but they must leave no trace in the final, gauge–invariant descriptors and pointer predictions.

Three tiers of "background". It is useful to sort what people call a "background" into three tiers:

(a) **Scaffolding (pure description).** Coordinates, gauge choices, tetrads, chart–dependent potentials, and similar bookkeeping devices. These *must* drop out of descriptors and predictions.

(b) **Effective backgrounds (slow sectors).** A slowly varying or stiff sector that is *treated as fixed* within a declared domain of validity—for example, a laboratory medium supplying an emergent metric g_{eff}, or a weak, slowly varying external control field. These are admissible when their status is declared, their calibration is supplied, and sensitivity to their variation is bounded.

(c) **Ontic stages (forbidden).** Postulated absolute space, time, or metric that is *not* part of the dynamics, not tied to devices, and yet is required to state physical predictions.

Principle 10.2: Background Independence as a Discipline.

State claims using descriptors that (i) are gauge/coordinate invariant, (ii) depend only on dynamical fields and declared boundary/control data, and (iii) remain stable under admissible changes of scaffolding and under small variations of any effective background across its domain of validity.

What it *forbids.*

- Promoting a chart, gauge, or foliation to ontology (reifying a convenient slicing as a "preferred frame").
- Building predictions that require a fixed metric not co–determined by the state (e.g. placing fields on rigid Minkowski space *and* letting their stress–energy have no back–reaction where such back–reaction would lie within the claimed regime).
- "Symmetry $\to$ number" shortcuts that extract absolute numerical values from invariance alone, without a dynamical mechanism or a calibration bridge (e.g. reading off a constant merely because a Lagrangian is scale–free).

What it *allows.*

- **Coordinates/gauges** chosen for convenience, provided final quantities are stated in invariants (curvature scalars, spectra, proper times/lengths, eigenvalues, relational ratios).
- **Effective media** that supply an emergent cone/metric g_{eff} (anisotropic optics, elastic media, superfluids). Here the "background" is a physical sector with proxies (densities, order parameters, stiffness), and predictions must include its calibration and a sensitivity bound.
- **External controls** (pump fields, trap potentials, rotation stages) treated explicitly as *device parameters* with units and ranges.
- **Symmetry reduction/ansätze** used to solve equations, so long as the discarded modes are shown not to affect the reported observables within the domain of validity.

Relational formulation. The safest expression of background independence is relational: predictions are phrased as relations among readings of calibrated devices carried by worldlines that can themselves be specified without hidden scaffolding. In practice this

means replacing "x, t at the lab" with, e.g., "phase difference between arms," "proper time between ticks," or "frequency ratio of two spectral lines measured by the same pointer chain."

Back–reaction and sensitivity. When using an *effective* background, declare a small parameter that justifies the treatment, and check it:

$$\epsilon_{\mathrm{BR}} \equiv \frac{\|\delta T_{\mu\nu}\|}{\|T^{\mathrm{bg}}_{\mu\nu}\|} + \left(\frac{\ell_{\mathrm{micro}}}{L_{\mathrm{bg}}}\right)^{\alpha} \ll 1,$$

where $\delta T_{\mu\nu}$ is the stress–energy of the excitations being modelled, $T^{\mathrm{bg}}_{\mu\nu}$ the supporting sector, ℓ_{micro} its micro–scale, L_{bg} the variation scale, and $\alpha > 0$ a model exponent. Predictions should vary by at most $O(\epsilon_{\mathrm{BR}})$ under admissible perturbations of the background proxies.

Worked contrasts.

- *GR vs. fixed stage.* Geodesic motion derived from a metric that satisfies dynamical equations (with specified boundary data) is admissible. The same geodesic law on a permanently rigid metric, while simultaneously claiming sensitivity to the traveller's stress–energy, is not.
- *QFT on curved spacetime (effective).* Treating $g_{\mu\nu}$ as a given field is acceptable as an *effective* description when back–reaction is bounded (declare ϵ_{BR}) and observables are stated in invariants of $g_{\mu\nu}$ and the detector model. Outside that bound, the description must be upgraded.
- *Emergent g_{eff} in media.* Birefringence cones, phonon cones, and similar g_{eff} are allowed when tied to measured proxies (density, stiffness, order parameters) and when the pointer predictions are insensitive to gauge and to small, calibrated drifts of those proxies.

Rule 10.2: How to write background–independent claims

R1. Say what is scaffolding. List coordinates/gauges and state that they are descriptive only.

R2. Name the effective background (if any). Provide its proxies, units, calibration, and a bound on ϵ_{BR}.

R3. State invariants. Rewrite the claim in invariant descriptors and pointer observables.

R4. Perturb and test. Show that the claim is stable under re–gauging and under small, admissible changes of the effective background across the declared domain

of validity.

What about constants of nature? Dimensionful standards (c, $\hbar$, k_B) are not backgrounds; they are *conversion factors* fixed by convention and measurement, entering only through units and calibrations. Dimensionless parameters (couplings) are part of the *model* and must be reported with their scale dependence and uncertainties.

Bottom line. Background independence here is a workflow: keep scaffolding out of ontology, demote slow sectors to calibrated controls with sensitivity bounds, and express everything in invariants with a device–facing bridge. Within this discipline theory becomes portable across coordinates, gauges, laboratories, and scales.

10.3 Measuring Law Strength

10.3.1 Across Scales

A *law* in this volume is a stable regularity: a mapping from descriptors (state, control, and boundary data) to pointer observables that survives changes of scale and description. By "scale" we include: spatial coarse–graining length ℓ, temporal bin size Δt, energy/frequency bandwidths, and averaging over ensembles. Strength is therefore not a single number but a *profile* that quantifies how far—and how uniformly—the regularity persists across these changes.

Set–up. Fix a scale window $W = [\ell_{\min}, L_{\max}]$ and a family of coarse–graining operators $\mathcal{C}_b$ (block average, low–pass filter, decimation) labelled by a factor $b > 1$. Let θ denote the effective parameters at a reference scale (e.g. couplings or transport coefficients). The coarse–grained dynamics induces a renormalization map $\mathcal{R}_b : \theta \mapsto \theta_b$ and a family of predictions $O_b = O(\theta_b)$ for the chosen pointer observable O. Uncertainties combine statistical and systematic budgets and are denoted σ_b.

Definition 10.3: Scale–stationarity index.

For an observable O, define the relative drift across a set of factors $\mathcal{B}$ by

$$\Delta_O \equiv \max_{b\in\mathcal{B}} \frac{|\, O_b - O_{b_0} \,|}{\sqrt{\sigma_b^2 + \sigma_{b_0}^2}} ,$$

where b_0 is a reference factor. The *stationarity index* is $\Sigma_O \equiv (1 + \Delta_O)^{-1} \in (0, 1]$; larger is stronger.

Definition 10.4: RG–stability radius.

Let $\|\cdot\|$ be a norm on parameter space. The RG drift over $\mathcal{B}$ is $D_{\mathrm{RG}} \equiv \max_{b\in\mathcal{B}} \|\theta_b - \theta_{b_0}\|/\varepsilon_\theta$, where ε_θ is the declared parameter uncertainty. The *stability radius* is $R_{\mathrm{RG}} \equiv (1 + D_{\mathrm{RG}})^{-1}$.

Definition 10.5: Extent of validity.

The *extent* E (in decades) is the largest interval of b for which the S/G/TQ/D gates are passed and Σ_O and R_{RG} exceed chosen thresholds:

$$E \;=\; \log_{10}\frac{L_{\max}}{\ell_{\min}} \quad \text{subject to} \quad \Sigma_O \geq \Sigma_\star,\; R_{\mathrm{RG}} \geq R_\star,\; \text{gates pass.}$$

Why these pieces? Σ_O measures *output robustness* (do the reported numbers survive coarse–graining?), while R_{RG} measures *model robustness* (do effective parameters wander beyond their error budgets?). The extent E captures how far the regularity reaches before gates or robustness fail.

Criterion 10.2: Composite law–strength score.

For a claim reported with one or more observables $O \in \mathcal{O}$, define

$$\mathcal{S}_{\mathrm{law}} \;\equiv\; \min_{O\in\mathcal{O}}\Big(\Sigma_O \wedge R_{\mathrm{RG}}\Big) \;\times\; \frac{E}{E_{\max}}\,,$$

where $E_{\max}$ is a community–chosen benchmark window for the field. Scores close to 1 indicate strong, portable laws; low scores point to scale–fragile patterns or narrow validity.

Across–scales workflow (practical test).

1. **Declare the window W and operators $\mathcal{C}_b$.** Justify choices physically (instrument bandwidths, sample sizes, relaxation).
2. **Compute/follow the flow.** Fit θ_b or derive $\mathcal{R}_b$; propagate to O_b with uncertainties σ_b.
3. **Report robustness.** Provide Σ_O, R_{RG}, and E with plots of $b \mapsto O_b$ and $b \mapsto \theta_b$.
4. **Gate check at each b.** No spectral pathologies, no superluminal signalling relative to g_{eff}, complete positivity for reduced maps, and non–negative dissipation at the reported timescale.

Rule 10.3: Law–strength audit (across scales)

L1. Parameter flow: show θ_b and the bound D_{RG}.

L2. Observable drift: show O_b with error bands and quote Δ_O and Σ_O.

L3. Extent: mark the largest contiguous b–interval where the gates pass; quote E.

L4. Sensitivity: perturb controls and effective background proxies within calibration budgets; confirm changes are $O(\varepsilon)$.

L5. Result: state $\mathcal{S}_{\mathrm{law}}$ and classify the claim (*provisional/strong/publication–ready*).

Interpretation. A high Σ_O with poor R_{RG} means the numbers look stable only because parameters are continually re–tuned (fragile "curve–fit" laws). A high R_{RG} with poor Σ_O indicates the model flows gently but the reported observable is not a true invariant of the flow—pick a better descriptor or a different pointer. A large extent E without gate failures is the most convincing sign of law–like behavior.

Worked example (schematic). Suppose a transport law predicts $J = \kappa(\theta)\,\nabla T$ for heat flux J. Under block–averaging by $b \in \{2, 4, 8\}$, the fitted conductivities κ_b vary by less than 5% within uncertainties ($R_{\mathrm{RG}} \approx 0.95$), and the coarse–grained fluxes J_b agree with J_{b_0} to within combined error ($\Sigma_J \approx 0.97$). The gates pass over two decades of ℓ, giving $E = 2$. With $E_{\max} = 3$ we obtain $\mathcal{S}_{\mathrm{law}} \approx 0.97 \times 0.95 \times \frac{2}{3} \approx 0.62$: robust and portable, but with a declared window.

Part VI

Beginnings and the Cosmic Narrative

This part treats a fraught question with operational discipline: how can a cosmic "beginning" be discussed by *physics* without importing metaphysical cargo? Our stance is simple: physics begins where a subject is *describable* in finite, gauge–invariant terms and connected, in principle, to a measurement pipeline. The story is told in state space rather than in absolute time; what counts are admissible descriptors and the dynamics they license.

Chapter 11

The Principle of Physical Beginning

This chapter states and motivates the *Principle of Physical Beginning*. It refuses "nothing" as a physical subject and replaces it with the earliest admissible *system* in which finite descriptors and causal dynamics can be posed. The aim is to separate physics from metaphysics while still enabling a testable narrative that can later be bridged to late–time observables.

11.1 Statement and Rationale

11.1.1 No "Nothing"

Principle 11.1: Principle of Physical Beginning.

A physical beginning is the earliest domain in state space where (i) at least one non–zero, finite, unit–bearing, gauge–invariant descriptor exists; and (ii) an admissible causal dynamics can be posed for that descriptor relative to an effective cone g_{eff}. The word "beginning" here names a *first describable system*, not a transition from "nothing".

Operational rationale.

1. **Measurability.** "Nothing" cannot be assigned units, a calibration, or an uncertainty budget; it fails the basic admissibility gates. Hence it is not a physical subject.

2. **Causality.** Dynamics requires a state space and an effective causal structure (an emergent cone g_{eff}). Without admissible descriptors, no update law can be posed.

3. **Gauge realism.** Claims about a beginning must be stated in gauge/coordinate invariants or accompanied by a clear disposal of gauge choices at the pointer. "Creation of space" in a specific chart is not an invariant statement.

4. **Stability at finite.** Any appeal to limiting ideals (e.g. continuum thermodynamic limits) must pass stability–at–finite audits. A narrative that needs literal infinities is outside the domain of validity.

Criterion 11.1: Admissibility checklist for "beginning" claims.

A beginning claim is *eligible for physics* only if it declares:

B1. Finite descriptors: the quantities used are finite and unit–bearing (no ∞ as a value).

B2. Gauge status: invariants are used, or gauge choices are flagged as descriptive and shown to cancel in device–facing outputs.

B3. Cone and dynamics: an effective causal cone g_{eff} and an admissible evolution law for the descriptors.

B4. Stability: a stated stability–at–finite budget for any idealizations.

B5. Bridge sketch: a plausible pathway from early descriptors to late–time pointer observables (even if indirect).

B6. Gate pass: S/G/TQ/D checks at the scales where the claim lives.

title=Language hygiene: "nothing" is not a subject

In this volume we *never* treat "nothing" as a dynamical antecedent, cause, or state. Phrases like "from nothing" are, at best, metaphors and must be rewritten as claims about the first admissible domain in state space, specified by finite descriptors and an effective dynamics.

State–space, not absolute time. The "first" here is not a moment on an absolute external clock. Following our internal–time discipline, the beginning is identified by the *minimal non–empty submanifold of state space* on which descriptors and dynamics are well posed. Temporal order is then read from admissible difference observables and relaxation scales, not imposed a priori.

What the principle forbids and allows. It *forbids* reifying coordinate artefacts (e.g. declaring a chart singularity to be a physical creation event), literal infinities as values, and dynamics without a state space. It *allows* effective media whose descriptors are crude at first, provided they are finite, gauge–stable, and come with an explicit plan for stability audits and empirical bridges.

Objection

"If you deny 'nothing', you have dodged the real question of why there is anything at all."

Reply

Physics does not answer "why existence?"; it answers "which admissible descriptors and dynamics, and with what empirical consequences?" The principle does not settle metaphysics; it fences physics so that its claims remain measurable, causal, and falsifiable.

Rule 11.1: Minimal reporting for beginning claims

Include (i) the list of finite descriptors with units and gauge status; (ii) a statement of the effective cone g_{eff} and evolution law; (iii) the stability–at–finite budget for all idealizations; (iv) a one–paragraph bridge to late–time observables; and (v) an S/G/TQ/D gate checklist.

The next section (*First Describable System*) proposes concrete templates for admissible seeds of g_{eff} and shows how rudimentary directional order, dissipation, and density–like descriptors can co–emerge to support causal propagation.

11.2 First Describable System

11.2.1 Seeds of g_{eff}

We now replace the slogan "no nothing" by a concrete, *first describable system.* The goal is not to guess the ultimate microphysics but to specify the minimal, finite set of descriptors that (i) admit a causal update law and (ii) determine an effective cone g_{eff} against which all device pipelines can be posed.

Definition 11.1: Minimal descriptor tuple (seed).

A first describable system is any medium that carries the tuple

$$\left(\rho,\ \eta,\ \phi,\ T,\ \chi^{\mu}\right),$$

with the following operational roles:

- ρ (*stiffness/density scale*): fixes characteristic speeds and susceptibilities.
- η (*dissipation/relaxation*): provides the arrow of time via non–negative entropy

production.

- ϕ (*phase/matter content*): supports energy shares and coherent distortions (order parameters).
- T (*internal clock*): enters only through differences/rates; absolute shifts are unobservable.
- χ^μ (*directional order*): a local orientation field that organizes transport; no universal preferred frame is imposed.

From seeds to a cone. A convenient schema that turns *scales* into a *cone* is

$$\begin{aligned} g_{\text{eff}}[\mu\nu](\rho,\chi) &= -\,c_\star^2(\rho)\,u_\mu u_\nu + \beta(\rho)\,h_{\mu\nu}, \\ u^\mu &= \frac{\chi^\mu}{\sqrt{\left|\eta_{\alpha\beta}\chi^\alpha\chi^\beta\right|}}, \quad u_\mu = \eta_{\mu\nu}u^\nu, \\ h_{\mu\nu} &= \eta_{\mu\nu} - u_\mu u_\nu. \end{aligned} \tag{11.1}$$

with positive functions α, β so that the characteristic speed $c_\star^2(\rho) = \beta(\rho)/\alpha(\rho)$ bounds signal/propagation cones. Equation (**??**) is not a commitment to a metric theory; it is a compact way to encode the *admissible set of trajectories* (cones) that subsequent dynamics must respect. Any alternative construction that yields the same cone class is equivalent for our purposes.

Minimal dynamical closure. Let ϱ_S denote the system state (to avoid overloading ρ). A generator consistent with the seeds is

$$\dot{\varrho}_S = -\,i\,[H(\phi,\rho),\varrho_S] \;+\; \sum_\alpha \gamma_\alpha(\eta,\rho)\Big(L_\alpha\varrho_S L_\alpha^\dagger - \tfrac{1}{2}\{L_\alpha^\dagger L_\alpha,\varrho_S\}\Big)\,, \tag{11.2}$$

with rates $\gamma_\alpha \geq 0$ (non–negative dissipation), pointer–facing observables built from energy shares of the L_α, and coarse–grained transport parameters constrained by the cone g_{eff}. Causality (no superluminal signalling *within* g_{eff}) and complete positivity are automatic if $\gamma_\alpha \geq 0$ and the L_α respect locality implied by (**??**).

Criterion 11.2: Seed admissibility.

A proposed seed $(\rho,\eta,\phi,T,\chi^\mu)$ is admissible when: (i) each entry is finite and unit–bearing; (ii) claims are expressed in gauge/coordinate invariants (or choices are flagged and eliminated at the pointer); (iii) g_{eff} derived from the seed is non–degenerate and closed under the evolution; and (iv) the stability–at–finite budget for any idealization is stated.

title=Device–level proxies at the beginning

Typical proxies: relaxation widths $\sim \eta$ (spectral lines, ring–down), birefringence or anisotropic scattering for χ^μ, wave speeds or elastic responses for ρ, phase–order parameters for ϕ, and rate/frequency differences for T. Early proxies may be indirect; that is acceptable if the bridge to later pointers is stated.

11.2.2 Prototype constructions

We record three minimal *templates*. They are examples of how the same seed tuple can instantiate g_{eff} and an admissible generator.

(A) Acoustic–like cone (hyperbolic transport). For a weakly compressible medium with flow v^i aligned with χ^μ and sound speed $c_\star(\rho)$,

$$ds^2 \propto -\left(c_\star^2 - v^2\right) d\tau^2 - 2\, v_i\, d\tau\, dx^i + \delta_{ij}\, dx^i dx^j, \tag{11.3}$$

which yields the same cone class as (**??**). Dissipation η converts high–frequency distortions into finite–width packets that respect the cone and stabilize pointer frequencies.

(B) Telegraph dynamics (finite–speed diffusion). A coarse–grained density q obeys

$$\tau(\eta)\, \partial_{TT} q + \partial_T q = D(\rho)\, \nabla^2 q, \tag{11.4}$$

with relaxation time $\tau(\eta) \geq 0$ and diffusivity $D(\rho) > 0$. Characteristics are bounded by $c_\star \sim \sqrt{D/\tau}$, again set by (ρ, η). In the $\tau \to 0$ limit, diffusion is recovered, but our admissibility requires working at finite τ (stability–at–finite).

(C) Open–system seed (Born–from–flows). Use (11.2) with jump operators that transport energy shares along χ^μ. Long–time stationary shares p_r read at the pointer obey $p_r = \mathrm{Tr}(E_r\, \varrho_*)$ with effects E_r proportional to calibrated energy–flow integrals (cf. the measurement part of the volume). The arrow of time rides on η via $\gamma_\alpha(\eta) \geq 0$.

Rule 11.2: Reporting template: first describable system

List the seed $(\rho, \eta, \phi, T, \chi^\mu)$ with units and gauge status; state your cone construction (e.g. (**??**) or (11.3)); give the evolution law (e.g. (11.2) or (11.4)); declare the stability–at–finite budget for all limits; and sketch the pointer bridge you intend to use at later times.

Failure modes (hard stops). (i) Any essential quantity $\to \infty$ as a value; (ii) cone degeneration or signalling outside g_{eff}; (iii) negative rates or ghost spectra; (iv) unstated/-failed stability–at–finite audits; (v) no device–facing bridge.

These templates are sufficient to continue the narrative without reifying "nothing": once a seed is declared and passes the gates, we can push its dynamics forward and design empirical bridges to the late Universe.

Chapter 12

Acceptance Checklist for a Beginning

This chapter turns the previous part into a one–page discipline for proposing a *physical beginning*. The goal is not to legislate metaphysics, but to make explicit the operational bar a proposal must clear before it counts as physics. We therefore state six acceptance conditions and pair each with concise tests and typical failure modes.

12.1 Six Conditions

Principle 12.1: Acceptance Principle (Beginnings).

A proposal qualifies as a *physical beginning* only if it is: (C1) describable in finite, unit-bearing terms; (C2) gauge-stable; (C3) explicit about its domain of validity; (C4) compatible with potential causality; (C5) equipped with a causal open-system dynamics candidate; and (C6) bridged to pointer observables with falsifiable signatures.

title=Summary of the six conditions

C1: Describability. All primitive descriptors are finite and unit-bearing (no bare use of ∞ as a value); their extraction procedures are at least conceptually specified.

C2: Gauge stability. Core claims survive admissible gauge/coordinate changes; when a gauge choice is used, it is flagged as descriptive and removed from observables.

C3: Declared DoV. The proposal states the scales, regimes, and approximations where it applies; end-of-validity indicators are identified *in advance*.

C4: Potential causality. The kinematic seeds admit an effective cone g_{eff} (or equivalent admissible transport structure) with no superluminal signalling relative to it.

C5: Causal dynamics candidate. There exists a generator consistent with the seeds and with S/G/TQ/D gates (e.g., a GKSL-type open evolution with complete positivity and non-negative dissipation on the reported timescales).

C6: Empirical bridge. Calibrated pointer observables are identified, with at least one discriminating, in-principle falsifiable signature at late times.

C1 — Describability. A beginning is physical only if the state of affairs is expressible by finite, unit-bearing descriptors (densities, rates, dimensionful constants) whose extraction procedures are at least sketched. When limits are used, a stability-at-finite audit must be supplied (tolerance and finite approximants).

C2 — Gauge stability. Statements are phrased in invariants (eigenvalues, curvature scalars, spectral data). Any dependence on coordinates or gauge is explicitly segregated and shown not to affect pointer outputs.

C3 — Declared domain of validity. The regime of applicability is stated in terms of scales (time, length, energy) and approximations (e.g., weak anisotropy, near-stationarity). Triggers for end-of-validity (blow-ups, loss of positivity, cone failure) are listed.

C4 — Potential causality. The seeds supply (or imply) an admissible transport structure—typically an effective cone g_{eff}—so that later dynamics and measurement pipelines can be embedded without acausal signalling.

C5 — Causal dynamics candidate. One exhibits at least a *candidate* open-system generator consistent with the seeds (e.g., Hamiltonian plus dissipators) that passes S/G/TQ/D: spectral safety for observables, causality within g_{eff}, complete positivity for reduced maps, and non-negative dissipation defining an arrow of time.

C6 — Empirical bridge. Concrete pointer observables and calibrations are named, together with a minimal late-time signature that would corroborate or falsify the narrative, distinct from competing accounts.

Reviewer's one-page checklist for a beginning

Report: (i) descriptors with units and finite values; (ii) invariants used; (iii) domain-of-validity and failure triggers; (iv) effective cone or transport structure; (v) a dynamics candidate and S/G/TQ/D audit; (vi) pointer observables, calibration sketch, and at least one discriminant.

Table 12.1: Six conditions and their quick operational tests.

Condition	**Quick test (to be reported)**	**If it fails**
C1 Describability	Units stated; no bare infinities; stability-at-finite tolerance given	Non-physical; move to mathematics only
C2 Gauge stability	Claim restated in invariants; gauge items flagged	Rephrase or drop ontology
C3 Declared DoV	Scales/regimes listed; end-of-validity triggers declared	Limit scope; add effective replacement
C4 Potential causality	Admissible cone/transport exists; no-signalling check	Reject or modify seeds
C5 Dynamics candidate	Generator passes S/G/TQ/D on reported timescales	Rebuild generator; report gate failure
C6 Empirical bridge	Named pointers; calibration sketch; discriminant given	Not physics until bridged

12.1.1 Failure Modes

F1 — Symbolic but non-physical. Descriptors involve ∞ as a value or lack units/stability-at-finite. *Action:* supply finite approximants with tolerance or demote the claim to a mathematical heuristic.

F2 — Gauge reification. A result changes under admissible coordinate/gauge transformations or uses a preferred chart as ontology. *Action:* rewrite in invariants; if impossible, classify as descriptive only.

F3 — Unstated scope. No scales or regimes are declared; predictions drift across coarse-graining. *Action:* declare a domain of validity; attach a limits audit plan.

F4 — Causal incoherence. Seeds do not admit an effective cone, or proposed transport enables signalling outside g_{eff}. *Action:* revise seeds or reject as non-physical.

F5 — Gate violations. Dynamics candidate fails spectral safety (ghosts/negative spectra), complete positivity on subsystems, or exhibits negative dissipation on the reported timescale. *Action:* rebuild the generator or restrict regime; report the failed gate explicitly.

F6 — No pointer bridge. No identified observables or calibrations; signatures indistinguishable from alternatives. *Action:* add a device-facing bridge and a concrete

discriminant; otherwise defer as speculative.

Rule 12.1: Hard stops (reject or rewrite)

Use of ∞ as a value (no stability-at-finite audit); gauge/coordinate artifacts treated as ontology; loss of complete positivity; negative dissipation on the reported timescale; absence of a domain-of-validity; missing pointer bridge.

12.2 Scope of Physics

This section draws a clean line between what *belongs to physics* in the beginnings narrative and what should be marked as *outside*. The point is not to police language, but to keep the claims we advance tied to descriptors, cones, generators, and pointer bridges that can in principle be checked.

Principle 12.2: Scope Principle.

A claim is *inside physics* iff it can be couched in finite, unit-bearing, gauge-stable descriptors, situated within an admissible transport structure (e.g., an effective cone g_{eff}), supplied with a domain of validity, and connected—directly or via a plausible path—to pointer observables. Otherwise it is *outside* and should be clearly marked as motivational or metaphysical.

Criterion 12.1: Inside/Outside test.

Let a candidate statement S be given. Then S is *inside* if it passes the six conditions of Section 12.1 (C1–C6). If any of C1–C3 fail, S is *outside*. If C4–C6 are pending but a concrete path-to-observables is supplied, classify S as *borderline (provisionally inside)* and attach milestones.

title=Green / Amber / Red tags (author discipline)

Green = inside: C1–C6 satisfied. **Amber** = borderline: C1–C3 satisfied and a time-bounded plan exists to meet C4–C6. **Red** = outside: one of C1–C3 fails (non-finite, non-invariant, or no declared DoV).

What stays inside. Typical inside items include: finite energy densities and rates at the putative beginning; statements about the emergence of an effective cone g_{eff} from directional order; a candidate GKSL generator consistent with S/G/TQ/D; and specific late-time discriminants (spectral slopes, anisotropy bounds, counting statistics).

What stays outside. Outside items include: creation "from nothing"; causes "before time"; teleological or value-laden purposes; bare use of infinity as a value; unrestricted multiverse assertions lacking a measure and discriminant; and global probability normalisations over an infinite ensemble with no stability-at-finite audit.

Table 12.2: Scope triage for beginnings claims.

Category	Description	Examples
Inside (Green)	Finite, invariant descriptors; g_{eff} present; DoV declared; bridge to pointers	Seed densities/rates; GKSL candidate with CP and non-negative dissipation; late-time spectral signature
Borderline (Amber)	Descriptors OK; cone plausible; path-to-observables listed with milestones	Toy effective medium with stated calibration plan; predicted imprint class but no instrument yet
Outside (Red)	Fails finite/invariant/DoV; or no conceivable pointer bridge	"Before time" causes; teleological narratives; bare ∞ as a value; measureless multiverse claims

Rule 12.2: Quarantine rule for metaphysical vocabulary

Use metaphysical or motivational terms only as *expository scaffolding*. Do not treat them as variables or inputs to equations; never let them determine a numerical prediction. When such terms appear, bracket them in prose and immediately translate to descriptors or drop them.

12.2.1 Inside/Outside

We adopt a short decision procedure for each beginning-related statement:

S1: Descriptors check (C1–C2). Are all primitives finite, unit-bearing, and invariant (or explicitly marked as choices)? If no, tag Red and remove from physics.

S2: Scope check (C3). Is the domain of validity declared with end-of-validity triggers? If no, tag Red until declared.

S3: Potential causality (C4). Do the seeds admit an effective cone g_{eff} (or equivalent admissible transport)? If not, revise the seeds; otherwise proceed.

S4: Dynamics candidate (C5). Is there at least a candidate generator passing

S/G/TQ/D on the reported timescales? If still under construction, tag Amber and record the gate tests to be run.

S5: Empirical bridge (C6). Are pointer observables and discriminants identified? If yes, Green; if not, Amber with a dated path-to-observables.

Border cases and how to handle them. *Typicality and selection effects* may be discussed inside physics only when a measure and sampling protocol are stated and a stability-at-finite test is reported. *Global initial conditions* are admissible when cast as constraints on invariant descriptors and when they feed a dynamical pipeline to pointers. Claims that depend essentially on uncheckable global structures without discriminants remain outside.

Minimal report to accompany a boundary claim

(i) Descriptors with units and invariance status.
(ii) DoV and end-of-validity triggers.
(iii) Cone or transport structure (existence of g_{eff}).
(iv) Dynamics candidate and which S/G/TQ/D checks are passed or pending.
(v) Pointer observables, calibration sketch, and at least one discriminant or a dated plan to obtain one.

12.3 Rules for Abstraction

Abstraction is indispensable: it compresses messy phenomena into tractable structures. But in a physics workflow abstraction is a *liability* unless it is disciplined by units, invariants, domains of validity, and a bridge to pointers. This section gathers the house rules that keep symbols honest and separates useful ideals from reified fictions. The spirit is simple: write claims the way an instrument would understand them.

Why rules at all? Because almost every crisis discussed in earlier parts can be traced to one of: (i) treating a limit as a value; (ii) mistaking a coordinate for an entity; (iii) omitting a domain of validity; or (iv) skipping the model $\to$ pointer bridge. The rules below make these failures explicit and prevent them at the source.

Rule 12.3: Abstraction rules (author discipline)

R1. Units first. No symbol enters without units or an explicit statement that it is dimensionless. Report the unit convention used.

R2. State invariants. Formulate claims in gauge/coordinate invariants; if a gauge

is chosen, flag it as scaffolding and show its disposal at the pointer.

R3. Declare DoV. Every formula carries a domain of validity (scales, approximations) and explicit end-of-validity triggers.

R4. Limits as tools. Any use of ∞, delta-idealizations, or exact symmetries must pass a stability-at-finite audit with a stated tolerance.

R5. Model $\to$ pointer bridge. Provide a concrete map from descriptors to device readouts and report calibration ingredients.

R6. Gate checks. Claims involving dynamics are accompanied by S/G/TQ/D checks on the stated timescale.

R7. Abstraction budget. Count your idealizations; each has a test plan and a date by which it will be stress-tested or retired.

From symbols to devices (the bridge). Let $\mathcal{L}$ be the effective generator (closed or open), $\mathcal{C}_\lambda$ a coarse–graining at resolution λ, and let $m_r(T)$ be the measured moment/rate in pointer channel r over an acquisition time T. The bridge is a calibrated map

$$\Pi_{\Delta t,\lambda} : \ (\rho_0, \mathcal{L}, \mathcal{C}_\lambda) \ \mapsto \ \{\, m_r(T) \,\}_{r=1}^{R}, \tag{12.1}$$

with sampling step Δt. Discipline requires: (i) *invariance*: Π yields the same numbers (within tolerance) under admissible re–gauging and coordinate changes; (ii) *stability at finite*: $m_r(T)$ stabilizes as T grows, with an uncertainty budget; (iii) DoV: Π is only claimed on the window where the device is linear and the hierarchy $\tau_B \ll \tau_{\text{mix}} \ll T$ holds.

12.3.1 Pointer Supremacy

Pointer supremacy is the pragmatic core of our approach: the pointer decides. When formal descriptions compete, the one that offers a cleaner, invariant, stable bridge to device readouts *wins* for physics; surplus structure that does not change pointers within the declared DoV is not to be reified.

Principle 12.3: Pointer Supremacy.

Between rival abstractions that agree on pointer statistics within the stated uncertainty and DoV, physics recognizes *equivalence*. Any surplus descriptive structure is scaffolding unless it produces a distinct, testable difference at the pointer.

This principle has two faces. *Conservative:* it forbids ontological weight to coordinate choices, gauge fixings, or ideal limits when they leave the pointer unchanged. *Aggressive:*

it pushes us to search for regimes where rival formalisms *do* diverge in pointer predictions (e.g., outside a Markovian window, under engineered non–CP noise, or with altered dissipation), thus turning metaphysical disagreements into empirical discriminants.

Criterion 12.2: Bridge adequacy test.

A claim invoking abstraction passes the supremacy test iff it supplies: (i) an explicit Π of the form (12.1) (or equivalent) with declared calibration steps; (ii) an invariance check (re–gauging/coordinate change); (iii) a stability–at–finite report (convergence of $m_r(T)$); and (iv) S/G/TQ/D checks on the timescale of T. Absence of any item demotes the claim to *abstract scaffolding*.

Worked contrasts (schematic). *Plane wave vs packet.* Using plane waves is admissible only after showing that the device–level averages for finite packets converge within tolerance to the plane–wave expressions on the declared window. Otherwise, the packet is the physics; the plane wave is a mnemonic. *Preferred foliation vs invariant report.* A calculation in a chosen foliation is harmless if the final report is cast in invariants (proper times, curvature scalars, spectral data) and shown insensitive to foliation changes; otherwise the foliation is being reified.

title=Abstraction budget (how to keep score)

List each idealization (infinite extent, exact symmetry, white noise, rigid background, delta source). For each: (a) where it enters the algebra, (b) the stability-at-finite test and tolerance, (c) which gate(s) it touches (e.g., positivity for white noise), and (d) the earliest regime where you will try to *remove* or *stress* it. This turns abstractions into testable assumptions rather than hidden commitments.

Ties and tie–breakers. If two models pass **??** and are empirically tied on the current window, we do *not* elevate hidden structure to ontology. Instead we rank by: (i) larger declared DoV; (ii) smaller abstraction budget; (iii) cleaner invariants; (iv) sharper empirical discriminants proposed for the next window. This keeps the practice forward–looking and device–anchored.

Bottom line. Abstraction earns its keep at the pointer. Rules **R1**–**R7** constrain how we write, while pointer supremacy tells us how to decide. The rest of the volume follows these rules: where we use an ideal, we say so and test it; where we must choose a description, we pick the one with the cleanest bridge; where rivals tie, we plan the experiment that breaks the tie.

Chapter 13

Early-to-Late Bridges

Claims about beginnings earn physical meaning only when we can sketch a path from *early descriptors* to *late-time pointers*. This chapter organizes such paths into reusable templates. We are not fixing a single cosmology; we are specifying how *any* admissible early state, written in finite, gauge-stable descriptors and with an effective cone g_{eff}, can leave calibrated footprints at later times.

title=Idea

A bridge is a workflow: early descriptors ⇒ effective dynamics ⇒ transport/relaxation ⇒ device-facing summaries. The bridge is accepted only if it declares its domain of validity, passes the S/G/TQ/D gates, and reports stability-at-finite for each idealization.

13.1 Bridging Scenarios

We catalogue scenarios that map seeds at the beginning into robust, pointer-level imprints. Each scenario states: (i) the *seed* (descriptor class), (ii) the *channel* (how it is transported through dynamics), and (iii) the resulting *signature* (what a device reads, with units and tolerances).

Set-up and notation. Let d_0 denote a vector of early descriptors (e.g., a stiffness/"density" proxy ρ_0, dissipation proxy η_0, and a directional-order proxy χ_0^μ). Let $\mathcal{U}_{t\leftarrow t_0}$ be the admissible evolution of descriptors to late time t, and let Π denote the calibrated model→pointer map (cf. Equation (12.1)). Then

$$\text{pointer readouts at time } t: \quad \mathbf{m}(t) \;=\; \Pi(\mathcal{U}_{t\leftarrow t_0}(d_0))\,, \tag{13.1}$$

with a declared domain of validity and uncertainty budget. The practical task is to choose scenarios for which $\mathbf{m}(t)$ features *stable* patterns that survive coarse-graining and gauge

changes.

13.1.1 Generic Signatures

We list signature families that are generic across wide classes of early seeds and dynamics. They are deliberately device-facing and admit falsification.

(S1) Monotone entropy production (arrow-of-time proxies). *Seed:* non-zero dissipation proxy $\eta_0 > 0$. *Channel:* open-system relaxation with admissible generator (e.g., GKSL-type). *Signature:* monotone growth of a calibrated entropy-production proxy $\dot{S}_{\text{prod}}$ or of variance-reduction rates in pointer channels, up to a plateau dictated by mixing. *Falsifier:* sustained negative dissipation at the reported timescale (fails D gate).

(S2) Directional fossils (anisotropy and birefringence bounds). *Seed:* small directional order $\chi_0^\mu \neq 0$. *Channel:* transport through an emergent g_{eff} with weak anisotropy. *Signature:* bounded, frequency-dependent anisotropy in calibrated spectra or time-of-flight differences; stable under re-gauging. *Falsifier:* dependence on coordinate choices (fails gauge realism) or lack of stability-at-finite under binning.

(S3) Scale-breaking knees (finite-range correlations). *Seed:* finite correlation scale L_0 in d_0. *Channel:* coarse-graining/transport that preserves a memory of L_0. *Signature:* "knees" in log–log spectral plots or in correlation functions at $k \sim 1/L_0$, surviving device bandwidth changes. *Falsifier:* knee drifts with binning alone (calibration artefact).

(S4) Non-Gaussian tails (limited mixing). *Seed:* non-linear couplings with incomplete mixing. *Channel:* weakly non-Markovian open dynamics. *Signature:* stable skewness/kurtosis excess in count statistics beyond Gaussian calibration, reported with stability-versus-T. *Falsifier:* disappearance under extended averaging consistent with pure device noise.

(S5) Cross-channel phase/energy shares. *Seed:* structured coupling among modes. *Channel:* mixing that settles to stable energy shares. *Signature:* late-time share vector $\mathbf{p}$ recovered from $\mathbf{m}(t)$ with total-variation $V(T)$ decaying as $T^{-1/2}$ on admissible windows; see Section 9.3. *Falsifier:* persistent plateaus in $V(T)$ within the declared window.

(S6) Echo asymmetry (arrow tied to openness). *Seed:* openness with finite memory. *Channel:* Loschmidt-echo manipulation. *Signature:* partial reversibility that scales with echo strength; residual irreversibility tracks dissipation proxies. *Falsifier:* echo has no effect on pointer statistics within uncertainty (undermines the openness-based arrow).

Rule 13.1: Minimal report for a bridging scenario

B1. Declare seeds d_0 (symbols, units, invariance status).

B2. Name the channel (evolution class and its gate checks).

B3. Specify the pointer (calibration, bandwidths, linearity range).

B4. State the signature (summary statistic, plot, and tolerance).

B5. Give the falsifier (what would make it wrong on the stated DoV).

Worked schematic (from S2). Early χ_0^μ induces a tiny birefringence in an emergent g_{eff}. Transport to late times yields a frequency-dependent phase delay $\Delta\phi(\omega)$. The pointer reads a calibrated slope $d\Delta\phi/d\log\omega$ that is: (i) invariant under re-gauging; (ii) stable under bandwidth changes within the window; and (iii) bounded by a function of the present-day proxies for the supporting medium. Reporting includes a stability-at-finite audit and an upper bound on χ_0^μ.

Criterion 13.1: Bridge acceptability (pass–fail).

A scenario is *accepted* for use in the narrative iff it declares its DoV and passes: (i) invariance check; (ii) stability-at-finite; (iii) S/G/TQ/D on the channel; (iv) a falsifier that is not purely linguistic. Otherwise it remains motivational and is bracketed as outside physics.

13.2 Toy Pipelines

We now exhibit compact, end–to–end *toy pipelines* that begin with early descriptors, pass through admissible dynamics, and finish at late–time pointer summaries. Each pipeline is deliberately simple, with declared domain of validity (DoV) and an explicit bridge of the form Π in Equations (12.1) and (13.1). The point is not realism; it is to make the logic inspectable and to show how seeds can be turned into falsifiable numbers.

Common ingredients. We write the system state as $\varrho(t)$ when quantum, or as a coarse distribution $f(x,t)$ when classical. Effective couplings are collected in θ with scale parameter ℓ (coarse–graining). An admissible evolution is specified either by a generator $\mathcal{L}_\ell$ (GKSL–type for open quantum, Fokker–Planck for classical) or by transport on an emergent cone described by $g_{\text{eff}}[\mu\nu]$. Pointer channels $r = 1, \ldots, R$ are fixed by a calibrated response and bandwidth.

13.2.1 Pipeline A: Diffusive seed → spectral knee

Seed (early descriptor). Finite correlation length L_0 in a scalar field with stiffness proxy ρ_0 and dissipation proxy $\eta_0 > 0$.

Channel (dynamics). Coarse diffusion on an emergent medium:

$$\partial_t f(x,t) \;=\; D(\rho_0,\eta_0)\,\nabla^2 f(x,t) \quad \text{with} \quad D>0, \tag{13.2}$$

valid on wavelengths $\lambda \gg \ell_{\text{micro}}$ and times $t \gg \tau_B$ (bath time).

Pointer and bridge. The device measures the power spectrum $S(k) = \langle |\tilde{f}(k)|^2 \rangle$ with a known window $W(k)$. The bridge Π reports the smoothed spectrum

$$\widehat{S}(k) \;=\; \int W(k-k')\,S(k')\,dk', \tag{13.3}$$

and the *knee* location $k_\star$ fitted from a broken–slope template.

Late–time signature. A stable knee $k_\star \sim 1/L_0$ persisting under device bandwidth changes within the DoV. *Falsifier:* $k_\star$ drifts under binning alone after deconvolution of W.

title=DoV and gates for Pipeline A

DoV: $\lambda \in [\ell_{\min}, L_{\max}]$ with $\ell_{\min} \gg \ell_{\text{micro}}$; linear response in the detector. Gates: spectral safety is trivial; causality is encoded by finite D and the cone of the medium; positivity holds for the Fokker–Planck semigroup; dissipation is non–negative.

13.2.2 Pipeline B: Directional seed → birefringence slope

Seed. Small directional order $\chi_0^\mu \neq 0$ that induces a weak anisotropy in an emergent metric–like descriptor $g_{\text{eff}}[\mu\nu]$.

Channel. Wave transport on g_{eff} with two polarization eigenmodes and phase velocities $v_\parallel, v_\perp$ differing by $\delta v(\omega)$ (linear regime).

Pointer and bridge. An interferometric pointer reports the phase difference

$$\Delta\phi(\omega) \;=\; \frac{\omega\,L}{v_\parallel(\omega)} \;-\; \frac{\omega\,L}{v_\perp(\omega)} \;\approx\; L\,\omega\,\frac{\delta v(\omega)}{v^2(\omega)}, \tag{13.4}$$

with L the calibrated baseline. The reported number is the slope $d\Delta\phi/d\log\omega$ over a window $\omega \in [\omega_1, \omega_2]$.

Late–time signature. A bounded, nonzero slope proportional to the seed anisotropy and *invariant* under re–gauging (same pointer, different description). *Falsifier:* slope depends on charting or disappears under bandwidth changes beyond the uncertainty budget.

title=DoV and gates for Pipeline B

DoV: weak anisotropy; linear optics; stable baseline. Gates: no negative spectra in the readout; causal propagation within the g_{eff} cone; complete positivity holds trivially at the classical field level; dissipation is either negligible on the window or reported separately as loss.

13.2.3 Pipeline C: Open quantum seed → energy–share vector

Seed. Coupled modes with coherent Hamiltonian $H(\theta)$ and weak openness to a bath supplying dissipators L_a with rates $\gamma_a(\eta_0)$.
Channel (GKSL).

$$\dot{\varrho} \;=\; -\frac{i}{\hbar}[H(\theta),\varrho] \;+\; \sum_a \gamma_a\Big(L_a\varrho L_a^\dagger - \tfrac{1}{2}\{L_a^\dagger L_a,\varrho\}\Big), \tag{13.5}$$

admissible on a window with $\tau_B \ll \tau_{\text{mix}}$.
Pointer and bridge. Channels r correspond to energy/counting outputs with calibrated POVM elements E_r. The measured rates $j_r(t) = \mathrm{Tr}(E_r\varrho(t))$ yield time–averaged shares

$$s_r(T) \;=\; \frac{1}{T}\int_0^T j_r(t)\,dt, \qquad \mathbf{s}(T) = (s_1,\dots,s_R) \;\xrightarrow[T\to\infty]{}\; \mathbf{p}. \tag{13.6}$$

Late–time signature. Stabilization of $\mathbf{s}(T)$ to a share vector $\mathbf{p}$ with $V(T) = \frac{1}{2}\sum_r |s_r(T) - p_r| \sim T^{-1/2}$. *Falsifier:* persistent plateaus in $V(T)$ on admissible windows, or gate failure (loss of CP, negative dissipation).

title=DoV and gates for Pipeline C

DoV: weak system–bath coupling; Markovian approximation on reported T; detector linearity. Gates: spectral safety of observables; causality inside the emergent cone; complete positivity by construction; non–negative entropy production.

13.2.4 Late-Time Pointers

Late–time readouts must be *device–facing summaries* with units and error budgets. We standardize four families:

P1. Spectral slopes and knees. Numbers extracted from $S(\omega)$ or $S(k)$ after window deconvolution; reported with stability versus binning and acquisition time.

P2. Phase/time delays. Differences like $\Delta\phi(\omega)$ or $\Delta t(\omega)$ with a baseline L; reported as slopes over stated windows and checked for gauge/representation invariance.

P3. Share vectors. $\mathbf{p}$ from $\mathbf{s}(T)$, reported with total–variation $V(T)$ versus T and Allan deviation; binning tests mandatory.

P4. Cumulant profiles. Skewness/kurtosis of count statistics, reported against calibrated Gaussian noise and with stability–at–finite audits.

Rule 13.2: Reporting template (late-time pointers)

For each pointer summary, include: (i) units and calibration map; (ii) DoV and gate checks; (iii) stability–at–finite plots (vs. T, binning, bandwidth); (iv) gauge/representation checks (same pointer, different description); (v) a falsifier specific to the chosen pipeline.

What makes a good toy pipeline. It should (a) start from finite, gauge–stable seeds; (b) use an admissible channel that passes S/G/TQ/D on the reported window; (c) land on a pointer summary whose numerical value is demonstrably stable under coarse–graining and gauge changes; and (d) include a concrete falsifier. These pipelines are the workhorses that carry the beginnings narrative from admissible seeds to measurable claims.

Part VII

Consistency Gates as Epistemic Virtues

This part codifies four *consistency gates* that every claim in the volume must pass on its declared domain of validity (DoV): **S**pectral safety, **G**eometric causality within an effective cone g_{eff}, **TQ** (complete positivity at the quantum level), and **D** (non–negative dissipation defining an arrow of time). Each gate is both a conceptual norm and a concrete test that can be run on models, numerics, and data products.

Chapter 14

Defining and Testing S/G/TQ/D

We give *operational* definitions of the gates and show how to test them. The aim is not to police formalisms but to protect device–facing claims from hidden pathologies. A model that clears the gates earns the right to appear in the narrative; a model that fails is either restricted in scope (DoV updated) or repaired.

14.1 Spectral Safety

The spectral gate Sgate demands that spectra relevant to reported observables and dynamics do not carry unphysical features such as negative norms, ghosts, or runaway eigenvalues on the stated timescale. The gate is checked at three layers: *kinematics* (state space and observables), *dynamics* (generators and propagators), and *responses* (correlation and power spectra read by devices).

Principle 14.1: Spectral Safety (Sgate).

Within the declared DoV: (i) physical observables A are self–adjoint with real spectra and calibrated ranges; (ii) quantum states are positive ($\varrho \geq 0$) and POVMs $E_r \geq 0$; (iii) generators produce contractive, spectrum–stable evolution (no eigenvalues with positive real part for dissipative semigroups; Hamiltonians bounded below); (iv) response spectra reported at the pointer are non–negative and consistent with positive–definite correlation kernels.

Kinematics (observables and states). Let $\mathcal{H}$ be the physical Hilbert space with positive inner product. For every reported observable A and calibration range $\mathcal{R}_A$:

$$A = A^\dagger, \qquad \mathrm{Spec}(A) \subset \mathbb{R}, \qquad \mathrm{Spec}(A) \cap \mathcal{R}_A \neq \emptyset. \tag{14.1}$$

Quantum states satisfy $\varrho \geq 0$, $\mathrm{Tr}\,\varrho = 1$; measurement elements $E_r \geq 0$, $\sum_r E_r = \mathbb{I}$. Any construction that requires an indefinite inner product or negative–norm subspace fails

Sgate.

Dynamics (propagation and generators). For unitary dynamics $U_t = e^{-iHt/\hbar}$ we require a self–adjoint H with

$$H \geq E_0 \mathbb{I} \quad \text{for some } E_0 \in \mathbb{R} \qquad \text{(energy bounded below).} \tag{14.2}$$

For open dynamics with a GKSL generator $\mathcal{L}$,

$$\frac{d}{dt}\varrho_t = \mathcal{L}(\varrho_t), \qquad e^{t\mathcal{L}} \text{ is CPTP for } t \geq 0, \tag{14.3}$$

the spectrum of $\mathcal{L}$ obeys $\operatorname{Re}\lambda \leq 0$ (contractive semigroup on trace norm). Numerically, one checks $\|e^{t\mathcal{L}}\|_{1\to 1} \leq 1$ on the window and absence of unstable eigenvalues.

Responses (what the device sees). Let $C_{XX}(\tau) = \langle X(\tau)X(0)\rangle$ be a stationary correlation. Positivity of C_{XX} as a kernel implies, by Bochner, that the power spectral density

$$S_{XX}(\omega) = \int_{-\infty}^{\infty} e^{-i\omega\tau} C_{XX}(\tau)\, d\tau \tag{14.4}$$

is non–negative. Any reported S_{XX} that dips below zero (after window deconvolution and uncertainty propagation) violates (Sgate). Likewise, the imaginary part of linear response $\chi(\omega)$ must satisfy standard positivity constraints compatible with dissipation (see **??**).

title=What (Sgate) forbids/permits

Forbids: wrong–sign kinetic terms (ghosts), non–self–adjoint observables used as if real, unbounded–above growth modes in a dissipative window, negative PSD after calibration. Permits: effective truncations that preserve self–adjointness and positivity; coarse models whose Lindbladians have $\operatorname{Re}\lambda \leq 0$ on the reported timescale.

14.1.1 Operator Tests

We collect practical tests that authors and reviewers can run on models, discretizations, and data products.

(T1) Self–adjointness and domains. State the operator domain and use a standard criterion:

$$H = H_0 + V, \quad H_0 \text{ self–adjoint and bounded below}, \quad \|V\psi\| \leq a\|H_0\psi\| + b\|\psi\|,\ a < 1 \Rightarrow H \text{ self–adjoint.} \tag{14.5}$$

On numerics, enforce symmetric discretizations and check that finite matrices approximate a Hermitian limit (eigenvalues real within tolerance).

(T2) Lower bound and ghost hunt. Verify (14.2) analytically (quadratic forms) or numerically (Rayleigh quotients). For field models, inspect the sign of kinetic terms; any negative coefficient in front of a squared time derivative signals a ghost.

(T3) CPTP and spectral bound for open maps. For a Lindblad form

$$\mathcal{L}(\varrho) = -\frac{i}{\hbar}[H, \varrho] + \sum_a \gamma_a \Big(L_a \varrho L_a^\dagger - \tfrac{1}{2}\{L_a^\dagger L_a, \varrho\} \Big), \tag{14.6}$$

check $\gamma_a \geq 0$ and compute the numerical spectrum of $\mathcal{L}$ on the truncated space used; confirm $\operatorname{Re}\lambda \leq 0$ on the DoV and $\|e^{t\mathcal{L}}\|_{1\to 1} \leq 1$ for sampled t.

(T4) PSD and kernel positivity. From data, build the empirical correlation $\widehat{C}_{XX}(\tau)$ and its PSD $\widehat{S}_{XX}(\omega)$. Report confidence bands after deconvolving device windows and subtracting known noise floors; require $\widehat{S}_{XX} \geq 0$ within uncertainty. As a cross–check, ensure that the Toeplitz matrices $[\,\widehat{C}_{XX}(t_i - t_j)\,]_{i,j}$ are positive semidefinite.

(T5) POVM and state checks. When reporting quantum readouts, reconstruct E_r and verify $E_r \geq 0$, $\sum_r E_r = \mathbb{I}$. Reconstructed states $\hat{\varrho}$ must satisfy $\hat{\varrho} \geq 0$, $\operatorname{Tr}\hat{\varrho} = 1$ within error; negative eigenvalues beyond uncertainty indicate a failed model/noisy inversion.

(T6) Discretization hygiene. For PDE models, use energy–stable schemes: discrete operators that mirror self–adjointness and dissipativity at the continuous level. Monitor the discrete energy or entropy; growth where the continuum predicts decay fails (Sgate) or (Dgate).

Rule 14.1: Spectral safety audit (author checklist)

S1. Declare domains for H, A, and the space on which $\mathcal{L}$ acts.

S2. Show $A = A^\dagger$ (or give the self–adjoint extension) and state the calibrated range.

S3. Prove/verify $H \geq E_0\mathbb{I}$ or give a transparent reason it holds.

S4. Give Lindblad form with $\gamma_a \geq 0$ and confirm $\operatorname{Re}\lambda \leq 0$ numerically on the truncated space used for plots.

S5. Report PSDs with non–negativity bands and Toeplitz positivity checks.

Typical failure modes and fixes. Wrong–sign kinetic terms (fix: correct the sign or restrict the DoV where higher–order terms repair the form); non–self–adjoint discretizations (fix: symmetrize and re–check spectra); unstable fitted Lindbladians with $\mathrm{Re}\,\lambda > 0$ (fix: enforce CP structure and refit under physical constraints); negative PSD after windowing (fix: recalibrate window, extend acquisition, or revise the model).

With (Sgate) in place, the next sections formalize (Ggate) (propagation inside an effective cone g_{eff} with no signalling outside it), (TQgate) (complete positivity on subsystems), and (Dgate) (non–negative dissipation and consistent arrows of time).

14.2 Causality within g_{eff}

14.2.1 Cones, No-Signalling

Causality in this volume is *relational and effective*: signals must propagate inside an *emergent cone* determined by device–facing descriptors, not by a pre–given stage. Concretely, suppose a metric–like descriptor $g_{\text{eff}}[\mu\nu](x)$ (or, more generally, a cone field) is inferred from calibrated proxies (e.g., stiffness, order, dissipation). Then admissible tangent vectors v^μ at x satisfy

$$g_{\text{eff}}[\mu\nu](x)\, v^\mu v^\nu \;\le\; 0, \qquad v^0 > 0, \tag{14.7}$$

and causal curves are those with tangent inside the cone a.e. For dispersive media, one replaces (14.7) by a *front–velocity* bound $c_\star(x)$ read from the high–frequency dispersion $\omega(k;x)$:

$$v_{\text{front}}(x) \;=\; \lim_{\omega\to\infty} \frac{\omega}{\|k\|} \;\le\; c_\star(x)\,, \tag{14.8}$$

with $c_\star$ itself a descriptor (finite, unit–bearing) tied to the same proxies. Superluminal *phase* or *group* velocities are not a violation so long as the front velocity obeys (14.8).

Operational no–signalling (device level). Pick two pointer channels, a *source* region R and a *target* region S, and define the causal travel time

$$\tau_\star(R\to S) \;:=\; \inf_{\gamma:R\to S} \int_\gamma \frac{ds}{c_\star(\gamma(s))}\,, \tag{14.9}$$

where γ ranges over curves tangent to the cone field. An intervention $\mathcal{I}$ supported in R must not change the marginal statistics at S before $\tau_\star$:

$$\mathrm{TV}\Big(P_S^{(\mathcal{I})}(t),\; P_S^{(\emptyset)}(t)\Big) \;=\; 0 \quad \text{for all } \; t < \tau_\star(R\to S), \tag{14.10}$$

where TV is total–variation distance between pointer distributions with/without the intervention. The statement is gauge/coordinate invariant because $\tau_\star$ is built from g_{eff} (or

$c_\star$) and calibrated lengths.

Field/response formulation. Let $G_{\text{ret}}(x, y)$ be the retarded Green function (or impulse response) for the linearized channel around the working point. Causality relative to g_{eff} is the support condition

$$G_{\text{ret}}(x, y) \;=\; 0 \quad \text{whenever } y \text{ lies outside the } g_{\text{eff}}\text{–past of } x. \tag{14.11}$$

Equivalently, in quantum settings with localized test functions f, g,

$$[\,A(f),\, B(g)\,] \;=\; 0 \quad \text{if } \operatorname{supp}(f) \text{ and } \operatorname{supp}(g) \text{ are } g_{\text{eff}}\text{–spacelike.} \tag{14.12}$$

For lattice/open systems, one uses Lieb–Robinson–type bounds: for operators supported on disjoint regions X, Y with distance $d(X, Y)$,

$$\|\,[\,A_X(t), B_Y\,]\,\| \;\le\; C\,\exp\!\Big(-\,[\,d(X,Y) - v_{\text{LR}}t\,]/\xi\Big)\,, \tag{14.13}$$

so the effective signal velocity is bounded by v_{LR}.

How to *measure* the cone. Excite a localized, calibrated source (impulse) and record arrival times across a detector array; fit isochron surfaces and regress $c_\star(x)$ (or fit a parametric $g_{\text{eff}}[\mu\nu]$) that best explains the isochrons. Cross–check with frequency–domain causality (analyticity): transfer functions $H(\omega)$ must be holomorphic in the upper half–plane and obey Kramers–Kronig relations; any violation flags a modelling/calibration error or a DoV breach.

Principle 14.2: Geometric causality (G).

On the declared DoV, influences propagate inside the effective cone; front velocities respect (14.8); retarded kernels satisfy (14.11); and device–level interventions pass (14.10). Dependence on charts/gauges is absent from the final statement.

Criterion 14.1: Gate G tests (pass–fail).

A claim involving propagation is admitted only if the author supplies:

G1. Cone descriptor: a calibrated $c_\star(x)$ or $g_{\text{eff}}[\mu\nu](x)$ with units and uncertainty.

G2. Kernel support: evidence that G_{ret} obeys (14.11) (analytic or numerical).

G3. Intervention test: an empirical or simulated check of (14.10) with randomized interventions.

G4. Bound check: a LR–style bound (14.13) (lattice/open) or a front–velocity

estimate (14.8) (continuum).

title=What does *not* violate causality?

Phase/group velocities exceeding $c_\star$ on a narrow band; apparent early ringing due to deconvolution artefacts; coordinate–dependent "superluminal" plots that disappear when expressed in invariants; parameter refits that change the *inferred* cone but not device–level no–signalling. These are not violations of G.

Rule 14.2: Causality audit (author checklist)

C1. Declare $c_\star$ or g_{eff} with units and calibration.

C2. Show retarded support or microcausality ((14.11)–(14.12)).

C3. Provide a no–signalling intervention plot satisfying (14.10).

C4. Quote a LR bound or front–velocity estimate with uncertainty.

C5. State the DoV and list conditions that would break G (e.g., strong nonlinearity, breakdown of the linear response fit).

The device-level no-signalling test is a pump–probe isochron check: before the front $t_f = L/v_f$ the pointer remains at noise level; cf. Figure 14.1.

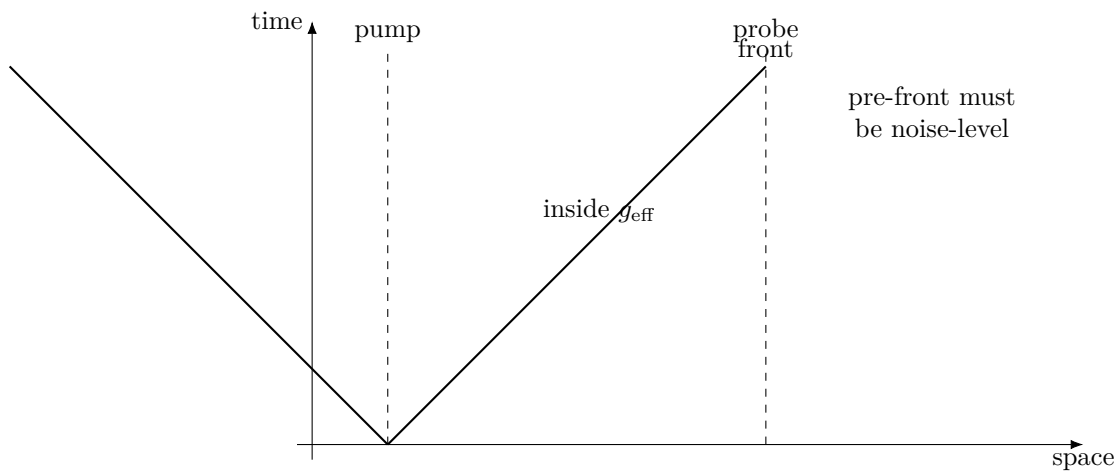

Figure 14.1: Pump–probe isochrons: no pointer change before the front time $t_f = L/v_f$ (no-signalling within g_{eff}).

Typical failure modes. Using group–velocity plots as "proof" of superluminal signalling; reporting coordinate times without invariant calibration; fitting transfer functions that violate analyticity; simulating open dynamics without a LR bound (spurious long–range couplings). Fixes: switch to front–velocity diagnostics, rewrite in invariants, enforce

causal filters, and add compact–support dissipators or finite–range couplings that admit a provable bound.

14.3 Complete Positivity

Complete positivity (CP) is the quantum version of the "no–negative-probabilities" rule for all *subsystems*. A physical process acting on a system S must map any valid state of S to a valid state *even when S is entangled with an uninvolved ancilla A of arbitrary size*. This is stricter than mere positivity and is the mathematically precise content of the TQ (quantum–positivity) consistency gate.

Definition 14.1: Completely positive, trace-preserving (CPTP) map

A linear map $\Phi : \mathcal{B}(\mathcal{H}_S) \to \mathcal{B}(\mathcal{H}_S)$ is *completely positive* if $\Phi \otimes \mathrm{id}_A$ is positive for every finite-dimensional ancilla A. It is a *quantum channel* when it is CP and trace-preserving (CPTP).

Principle 14.3: TQ (quantum positivity)

Reduced dynamics and reported instruments must be CPTP on their declared domain of validity. Where CP fails (e.g., due to initial system–environment correlations), either (i) enlarge the system so the map becomes CPTP, or (ii) restrict the preparation class and model the process as a higher-order object (superchannel/process-tensor) rather than as a state map.

Three equivalent CP characterizations (toolkit). For finite dimensions, the following are equivalent and give practical tests:

(a) **Kraus form.** $\Phi(\rho) = \sum_k K_k \rho K_k^\dagger$, with $\sum_k K_k^\dagger K_k = \mathbb{I}$ (trace-preserving case).

(b) **Stinespring dilation.** $\Phi(\rho) = \mathrm{Tr}_E\Big[U(\rho \otimes |0\rangle\langle 0|)U^\dagger\Big]$ for some unitary U and environment E.

(c) **Choi positivity.** The Choi operator $J_\Phi = (\Phi \otimes \mathrm{id}_A)(|\Omega\rangle\langle\Omega|) \geq 0$, where $|\Omega\rangle = \sum_{i=1}^d |i\rangle \otimes |i\rangle / \sqrt{d}$.

Operationally, (c) is the basis of ancilla/entanglement tests; (a) is ideal for analysis and simulation; (b) connects to Hamiltonian-and-bath models.

Time–local generators and CP. A time–local master equation

$$\dot{\rho}_S(t) = -i[H(t), \rho_S(t)] + \sum_\alpha \gamma_\alpha(t)\Big(L_\alpha(t)\rho_S(t)L_\alpha^\dagger(t) - \tfrac{1}{2}\{L_\alpha^\dagger(t)L_\alpha(t), \rho_S(t)\}\Big) \tag{14.14}$$

generates a CP map for each time interval provided the *Kossakowski matrix* is positive semidefinite, equivalently the rates $\gamma_\alpha(t) \geq 0$ almost everywhere. This is the GKSL structure; violation (negative rates) is a red flag for both TQ and D.

title=Common failure modes (what TQ catches)

- Using merely *positive* (not CP) maps, e.g. matrix transposition or phenomenological "non–CP" master equations.
- Truncating time–convolutionless expansions so that the effective Kossakowski matrix acquires negative eigenvalues.
- Ignoring *initial system–environment correlations*: reduced dynamics need not be CP unless the assignment from system preparations to joint states is CP.
- Reporting instruments that are not CP on *coarse-grained bins* (e.g., deconvolution artifacts yielding negative POVM elements).

Criterion 14.2: TQ admission test (the map level)

A reported dynamical map $\Phi_{t:s}$ or measurement instrument $\{\mathcal{I}_r\}$ is admissible only if at least one of the following is supplied (with units/DoV):

C1. A Kraus/Stinespring representation on the declared DoV.

C2. A Choi operator estimate $J_\Phi \geq 0$ (tomography or a certified lower bound).

C3. A GKSL generator with positive semidefinite Kossakowski matrix throughout the applicable timescale.

Violations must be flagged and explained (e.g., restricted preparation class, non-Markovian revivals, or model incompleteness).

14.3.1 Subsystem Consistency

Subsystem consistency demands that local modeling choices do not produce contradictions when the system is viewed as a part of a larger whole. In practice this has three layers:

(i) Ancilla–safety. For any ancilla A uninvolved in the dynamics, $(\Phi \otimes \mathrm{id}_A)$ must map *all* joint states to states. This is precisely CP. A convenient experimental surrogate is the *Choi ancilla test*: prepare a maximally entangled SA pair (or a certified entanglement proxy), act with Φ on S, then reconstruct a witness that lower-bounds $\lambda_{\min}(J_\Phi)$. A negative bound indicates failure.

(ii) Composition–safety. Independent devices should compose to a channel on the composite system. If two reported channels Φ^A, Φ^B act on subsystems simultaneously, then $\Phi^{AB} = \Phi^A \otimes \Phi^B$ must be CPTP on the stated DoV. Deconvolutions and crosstalk corrections must be shown not to introduce non-CP effective maps at the composite level.

(iii) Reduction–safety (initial correlations). When S is initially correlated with an environment E, the reduced map on S can fail to be CP. Two CP-safe routes:

(i) **Enlarge** the modeled system to $S' \supseteq S$ so that the joint dynamics is CPTP and the final readout is a partial trace on S'.

(ii) **Restrict** the preparation class to those generated by a CP assignment Λ ("preparation procedures"), then the composed map $\Phi \circ \Lambda$ is CP on that class. Report the restriction explicitly.

Proposition 14.1. *Data-processing sanity checksprop:dpi If Φ is CPTP, then for any contractive information measure $D(\cdot\|\cdot)$ (e.g. quantum relative entropy, trace distance), $D(\Phi(\rho)\|\Phi(\sigma)) \leq D(\rho\|\sigma)$. Empirical violations (within uncertainties) signal either non-CP processing in the pipeline or unreported DoV breaches.*

Rule 14.3: Subsystem–consistency audit (author checklist)

SC1. Declare whether ancillas may be entangled with S in the experiment/theory; if yes, provide a CP proof (Kraus/Stinespring or Choi).

SC2. If initial SE correlations exist, either enlarge S or state the CP assignment map/restricted preparation class.

SC3. For parallel devices, certify that crosstalk corrections keep $\Phi^A \otimes \Phi^B$ CPTP on the composite DoV.

SC4. For time–local models, show that all reported rates are non-negative on the timescale used (shared with D).

SC5. Provide at least one ancilla–based or Choi–based witness (or a tomography reference) that supports CP within uncertainty.

title=Reporting template (CP-safe)

Map: $\Phi_{t:s}$ on S; **DoV:** (states, energies, times); **Representation:** Kraus $\{K_k\}$ / Choi J_Φ / GKSL rates $\{\gamma_\alpha\}$; **Ancilla policy:** allowed/forbidden; if allowed, CP witness result; **Preparation class:** unrestricted / restricted by CP assignment; **Coarse-graining:** bin width, deconvolution method; **Gate status:** TQ satisfied (evidence), D (rate signs), S (no negatives), G (no-signalling).

Typical pitfalls and remedies.

- *Non-CP "fixes" to data* (e.g. naive background subtraction making POVM elements negative). **Remedy:** constrain estimation to the CPTP cone (maximum-likelihood with CP constraints).
- *Negative rates in fitted master equations.* **Remedy:** refit within GKSL structure or report a non-Markovian model (and move the claim outside TQ).
- *Hidden initial correlations.* **Remedy:** state a CP assignment (with calibration) or expand the modeled system.

14.4 Non-Negative Dissipation

Dissipation is where the arrow of time enters our modelling in a device-facing way. At the level of reduced open dynamics, the D demands that reported generators and coarse-grained balances do not allow net entropy to *decrease* on the instrument's timescale. This is stronger than mere energy bookkeeping: it is a constraint on the *sign of irreversibility* compatible with calibration and uncertainty.

Definition 14.2: Entropy balance and currents (lab convention)

Let $\rho_S(t)$ be the system state, $H(t)$ its (effective) Hamiltonian, and $\dot{\rho}_S(t) = \mathcal{L}_t[\rho_S(t)]$ a reported time-local generator. We split energy change as

$$\frac{d}{dt}\,\mathrm{Tr}[H\rho_S] \;=\; \underbrace{\mathrm{Tr}[H\,\mathcal{L}_t(\rho_S)]}_{\dot{Q}\ \text{(heat into } S)} \;+\; \underbrace{\mathrm{Tr}[(\partial_t H)\rho_S]}_{\dot{W}\ \text{(power/work)}}\,. \tag{14.15}$$

The von Neumann entropy rate is $\dot{S}_S(t) = -\,\mathrm{Tr}[\dot{\rho}_S \ln \rho_S]$. Given a set of labelled environments $\{E_\alpha\}$ (baths, measurement ports), we decompose the dissipator and heat currents as $\mathcal{L}_t = \sum_\alpha \mathcal{L}_t^{(\alpha)}$ and $\dot{Q}_\alpha = \mathrm{Tr}\Big[H\,\mathcal{L}_t^{(\alpha)}(\rho_S)\Big]$.

Principle 14.4: D: Non-negative dissipation on the operating timescale

Reported dynamics and instruments must be consistent with a non-negative entropy production rate (EPR) on the declared domain of validity and temporal resolution T_{res}:

$$\dot{\Sigma}(t) \;\equiv\; \dot{S}_S(t) \;-\; \sum_\alpha \beta_\alpha(t)\,\dot{Q}_\alpha(t) \;\geq\; 0 \quad \text{(within uncertainty).} \tag{14.16}$$

Here β_α are environment inverses of temperature (or generalized affinities when relevant). If instantaneous oscillations appear below T_{res}, the coarse-grained production

$\Sigma_{[t,t+T_{\mathrm{res}}]} = \int_t^{t+T_{\mathrm{res}}} \dot{\Sigma}(\tau)\, d\tau$ must be non-negative.

Spohn's inequality for GKSL semigroups. When $\mathcal{L}$ is time-independent of Gorini–Kossakowski–Sudarshan–Lindblad form and admits a stationary state ρ_* (e.g., thermal), Spohn's theorem gives

$$\frac{d}{dt} D(\rho_S(t)\|\rho_*) \;=\; \mathrm{Tr}\Big[\mathcal{L}(\rho_S)\,(\ln \rho_S - \ln \rho_*)\Big] \;\le\; 0, \tag{14.17}$$

i.e. the relative entropy to equilibrium is non-increasing. For a thermal fixed point with $\rho_* \propto e^{-\beta H}$ this yields (14.16) with a single bath, ensuring $\dot{\Sigma} \ge 0$.

Local detailed balance (multi-bath models). In standard weak-coupling constructions with jump operators $L_{\alpha,\omega}$ and rates obeying the KMS relation $\gamma_\alpha(-\omega) = e^{-\beta_\alpha \omega}\gamma_\alpha(\omega)$, one obtains the operational form

$$\dot{\Sigma}(t) \;=\; \dot{S}_S(t) \;-\; \sum_\alpha \beta_\alpha\, \dot{Q}_\alpha(t) \;\ge\; 0, \qquad \dot{Q}_\alpha(t) = \sum_\omega \omega\, J_{\alpha,\omega}(t), \tag{14.18}$$

with $J_{\alpha,\omega}$ the net quanta currents. Violation indicates either a modelling error (rates not obeying detailed balance on the declared DoV) or that the coarse-graining/Markov assumptions have been exceeded.

title=What D does *not* require

D is a statement about *reported* reduced dynamics on an instrument's resolution. Transient negativity at sub-resolution, or apparent violations arising from mis-assigned work/heat partitions, are not failures once the DoV and calibration are stated and the coarse-grained production is non-negative.

14.4.1 Entropy Production

Entropy production (EP) is the quantitative footprint of irreversibility, and it can be certified at three complementary levels:

(1) Generator level (model audit). Provide either:

- a GKSL generator with a positive semidefinite Kossakowski matrix and a stationary state ρ_*; use (14.17) to conclude $\dot{\Sigma} \ge 0$; or
- a multi-bath construction satisfying local detailed balance; report β_α and the decomposition (14.18).

(2) Balance level (energy–entropy accounting). Using (14.15), reconstruct $\dot{Q}_\alpha$ and $\dot{W}$ from measured powers and temperatures, then compute $\dot{\Sigma}$ via (14.16). Uncertainty budgets (sensor noise, calibration, binning) must accompany any claim of sign.

(3) Trajectory level (fluctuation theorems). For repeated runs with trajectory inference (e.g., quantum jump records), the total stochastic EP obeys the integral fluctuation theorem

$$\left\langle e^{-\Delta s_{\mathrm{tot}}} \right\rangle = 1 \qquad \Rightarrow \qquad \langle \Delta s_{\mathrm{tot}} \rangle = \Sigma \geq 0, \tag{14.19}$$

providing an empirical witness of irreversibility consistent with D.

Criterion 14.3: D admission test (reporting)

A claim passes D when the author provides: (i) a clear heat/work split (14.15) with units and DoV; (ii) either a GKSL or detailed-balance certification ensuring $\dot{\Sigma} \geq 0$, *or* a trajectory-based FT check (14.19); (iii) an uncertainty analysis showing non-negativity within resolution T_{res}.

Rule 14.4: Dissipation audit (author checklist)

D1. Declare baths/ports and their β_α (or generalized affinities), and the instrument resolution T_{res}.

D2. State the energy balance (14.15) and how $\dot{Q}_\alpha$ and $\dot{W}$ are obtained from data.

D3. Give either a GKSL proof (rates ≥ 0, fixed point) or KMS/local-detailed-balance relations; otherwise use a trajectory FT witness.

D4. Report $\dot{\Sigma}$ and the coarse-grained $\Sigma_{[t,t+T_{\mathrm{res}}]}$ with uncertainties.

D5. Flag and explain any apparent negative segments (DoV breach, non-Markovian backflow, or mis-calibration) and show that coarse-grained EP is non-negative.

Typical pitfalls and cures.

- *Negative fitted rates* in a time-local master equation. **Cure:** refit within GKSL structure or enlarge the model to include memory; restrict claims outside D if necessary.
- *Heat/work mis-assignment* when $H(t)$ is driven. **Cure:** use (14.15); power is $\mathrm{Tr}[(\partial_t H)\rho]$, not absorbed heat.
- *Hidden leakage channels* (unmodelled baths) that flip the EP sign. **Cure:** add ports α with calibrated β_α or restrict the DoV.

Chapter 15

Gate-Based Audits

The aim of a gate-based audit is to surface, in a uniform way, whether a given claim or model obeys the four consistency gates introduced in Part VII: S (spectral safety), G (causality within g_{eff}), TQ (complete positivity on subsystems under the stated reductions), and D (non-negative dissipation on the reported timescale). Each audit is device-facing: it couples a brief mathematical check to an operational reading plan and an explicit domain of validity (DoV).

15.1 Three Mini Case Studies

15.1.1 GR/QFT/Media

We illustrate the audit on three familiar contexts. The point is not to re-prove theorems, but to show how *the same* checklist clarifies what is claimed, what is measured, and where the limits lie.

Case A: GR horizon vs singularity (Schwarzschild)

Setup. Consider the Schwarzschild solution with mass M. An often–confused contrast is between the event horizon at $r = 2M$ and the central singularity at $r = 0$. A curvature invariant that diagnoses intrinsic blow-up is the Kretschmann scalar

$$K(r) = R_{abcd}R^{abcd} = \frac{48\,M^2}{r^6}. \tag{15.1}$$

Rule 15.1: Gate audit (GR horizon vs singularity)

1. **S.** Observables formed from curvature invariants are finite at $r = 2M$ but diverge at $r \to 0$ by (15.1). *Status:* horizon passes; central point fails (end-of-validity flag).

2. **G.** Causal cones of g_{eff} follow the metric light cones. Horizon is a global causal boundary, not a local superluminal effect. *Status:* pass.
3. **TQ.** Not directly applicable in a closed classical GR solution; for reduced quantum fields on this background, CP must be checked at the level of the reduced dynamics (outside scope here). *Status:* neutral.
4. **D.** Classical vacuum solution has no device-level entropy production; any dissipative claim must come from matter channels or coarse-grained quantum fields. *Status:* neutral.

Outcome and DoV. Claims that reify the horizon as a curvature blow-up *fail* S and should be rephrased in invariant terms. The singularity at $r = 0$ is an *end-of-validity* signal for the classical description, triggering the replacement workflow of **??**.

Case B: QFT ultraviolet divergences and renormalization

Setup. Take scalar φ^4 theory in four dimensions. A one-loop self-energy exhibits a UV divergence, e.g.

$$\Sigma(p) \;\propto\; \lambda \int^{\Lambda} \frac{d^4k}{(2\pi)^4} \frac{1}{(k^2+m^2)((k+p)^2+m^2)}. \tag{15.2}$$

Renormalization defines finite, scale-dependent couplings that feed the RG flow.

Rule 15.2: Gate audit (QFT UV and RG)

1. **S.** Källén–Lehmann representation enforces a non-negative spectral density for physical fields; renormalized propagators must avoid negative-norm ghosts within the DoV. *Status:* pass if the spectral density is non-negative and counterterms preserve unitarity.
2. **G.** Microcausality requires vanishing commutators at spacelike separation; the renormalized theory must retain this support property. *Status:* pass if locality is preserved.
3. **TQ.** When coarse-graining to an *open* effective sector (integrating out modes), any Markovian reduction used for predictions must be CP. A purely closed scattering claim is neutral. *Status:* conditional pass (declare the reduction and show CP).
4. **D.** Closed QFT has zero entropy production; coarse-grained kinetic descriptions must show non-negative production on the stated timescale. *Status:* neutral for S-matrix; pass for kinetic models with positive collision entropy.

Outcome and DoV. The divergence in (15.2) is an instance of "infinity as a limit". After renormalization, claims are confined to energy windows where the effective couplings keep S and G intact; outside that window the model's DoV ends.

Case C: Dispersive media and "superluminal" group velocity

Setup. In a passive, causal, linear medium the susceptibility $\chi(\omega)$ is analytic in the upper half-plane and obeys Kramers–Kronig relations. Narrow-band pulses can display $v_g > c$ while the front velocity remains bounded by $c_\star$ (the g_{eff} cone).

Rule 15.3: Gate audit (dispersive media)

1. **S.** Passivity implies a positive-real response and non-negative absorption: $\operatorname{Im}\chi(\omega) \geq 0$ for $\omega > 0$. No negative-energy modes. *Status:* pass with a calibrated response fit.
2. **G.** No-signalling is guaranteed by analyticity (fronts propagate at $c_\star$). Apparent $v_g > c_\star$ arises from pulse reshaping, not cone violation. *Status:* pass when the no-signalling intervention plot is shown.
3. **TQ.** Quantum models of the medium (bath couplings) must include fluctuation–dissipation compatible noise to keep reduced dynamics CP. *Status:* conditional pass (show CP of the reduced map).
4. **D.** Passivity yields non-negative dissipation on the operating band; fitted negative loss indicates miscalibration or DoV breach. *Status:* pass with uncertainty budget.

Outcome and DoV. Claims of "superluminal signalling" fail G once the front velocity and intervention protocol are reported. What remains is a shape effect consistent with the gates on the declared band and device resolution.

These small audits demonstrate how the same four-gate lens suppresses reification (coordinate artefacts in GR), domesticates infinities (QFT UV), and polices causal language in effective media. In Section 15.2 we turn this into a one-page template authors can attach to any substantive claim.

15.2 Writing the Gate Report

A gate report is a one–page, device–facing audit that ties a claim to: (i) gauge–invariant, finite descriptors; (ii) a model→measurement pipeline; and (iii) the four gates: **S** (spectral safety), **G** (causality within g_{eff}), **TQ** (complete positivity under the stated reductions), and **D** (non–negative dissipation on the reported timescale). It also declares the domain

of validity (DoV) and the uncertainty budget so that other readers can reproduce the check.

How to use this section. Authors attach an *Author Gate Card* to each substantive statement or figure; reviewers may use the *Reviewer Card* to request any missing items. Both cards are meant to be concise and uniform across theory, modelling, and experiment.

15.2.1 Templates

title=Author Gate Card (one page)

Claim (one sentence in invariants with units).
Descriptors and units: list symbols, units, and gauge status (invariant vs choice).
Pipeline: model → effective dynamics (closed/open) → pointer observables; calibration reference.
DoV: regime/scales, approximations, cutoffs; conditions that would end validity.
Uncertainty budget: statistical, calibration, estimator bias; stability-at-finite check.
S gate (spectral safety): operator/spectrum test; absence of negative modes/ghosts.
G gate (causality within g_{eff}): cone or no–signalling intervention plot; micro/retarded support.
TQ gate (complete positivity): state the reduction (partial trace/coarse graining); show CP of the reduced map or cite a theorem that applies.
D gate (non–negative dissipation): rate signs/entropy production on the reported timescale.
Outcome: Provisional / Strong / Publication–ready (choose one) and give a one–line justification.

title=Reviewer Gate Card (request/verify)

1. Restate the claim in invariant, unit–bearing form.
2. Ask for the DoV and the stability–at–finite evidence.
3. Check S: spectrum/positivity tests attached?
4. Check G: cone or no–signalling intervention plot attached?
5. Check TQ: is the stated reduction CP (theorem or calculation)?
6. Check D: are reported rates/entropy production non–negative on the stated window?
7. Calibration and units: is there a pointer–level bridge?
8. Verdict and *missing items list.*

Rule 15.4: Minimal Gate Form (fillable)

Claim: .....
Descriptors (units / gauge status): .....
Pipeline: .....
DoV (regimes/scales): .....
Uncertainty & stability-at-finite: .....
S: operator/spectrum check: .....
G: cone or no–signalling check: .....
TQ: CP of reduced map: .....
D: rate/entropy sign on timescale: .....
Outcome: Provisional | Strong | Publication–ready

Example (filled skeleton; theory→pointer). The following toy example shows the level of specificity expected. Replace placeholders with your quantities.

Rule 15.5: Example Gate Report (toy)

Claim. In the band $[\omega_1, \omega_2]$ the invariant damping rate $\Gamma(\rho)$ of mode k_* is non–negative and satisfies $\Gamma(\rho) \leq \Gamma_{\max}$ (units s^{-1}).
Descriptors. ρ (stiffness, Pa), η (dissipation, Pa·s), ϕ (phase), χ^μ (directional order); all reported as gauge–invariant combinations; units stated.
Pipeline. Effective linear response $\Rightarrow$ open dynamics with GKSL generator; pointer observable is the calibrated energy share in channel r; calibration via reference sample C.
DoV. Strain $< 10^{-3}$, temperature 300 ± 1 K, band $[\omega_1, \omega_2]$; linearity verified; no strong non–Markovian revivals on $T \leq 10$ ms.
Uncertainty. Statistical $\pm 0.3\ \mathrm{s}^{-1}$; calibration $\pm 0.2\ \mathrm{s}^{-1}$; estimator bias $< 0.1\ \mathrm{s}^{-1}$; stability-at-finite checked by windowing $N = 2^m$ with tolerance $\varepsilon = 0.02$.
S. Fitted generator has non–negative spectrum for observables (no ghosts); residuals within noise.
G. No–signalling intervention plot (pump–probe) shows fronts within the inferred g_{eff} cone.
TQ. Reduced map (system vs bath) is CP; proof via standard GKSL structure with positive rates.
D. Entropy production $\dot{S} \geq 0$ on $T \leq 10$ ms; negative estimates disappear within uncertainty.
Outcome. *Strong*: baseline + quantitative uncertainty budget + reproducible script.

Tabular rubric (pass/fail per gate). Use the rubric to summarize gate status succinctly.

Table 15.1: Gate rubric: what counts as a pass/fail and common fixes.

Gate	Pass (evidence)	Fail (symptom) / Fix
S	Positive spectra; no negative–norm modes; Källén–Lehmann–type density non–negative.	Ghost/negative mode ⇒ fix model or restrict DoV.
G	Retarded/microcausal support; device–level no–signalling plot within g_{eff}.	"Superluminal" from coordinate artifacts or deconvolution ⇒ re–express in invariants; show front plots.
TQ	CP of reduced map under stated coarse graining.	Non–CP fit ⇒ add noise channels or revise reduction.
D	Non–negative dissipation/entropy on reported window.	Negative rate from bias/under-modelled feedback ⇒ recalibrate or shorten DoV.

Strength tiers (beyond baseline). Baseline = all four gates reported with DoV and units. *Strong* = baseline + uncertainty budget and robustness checks. *Publication–ready* = strong + falsifiable empirical signature (scale/pattern) and a reproducible script/pseudocode for the gate tests.

Part VIII

Virtues and Limits of Mathematics

Mathematics is the language we use to compress, relate, and project. In this part we ask what mathematics can certify for physics *without* overstepping the boundaries set by measurability, causality, domains of validity (DoV), and the model→pointer bridge. Our stance is friendly but firm: proofs are powerful inside a stated structure; they do not by themselves enlarge a model's DoV or create a physical subject.

Chapter 16

Proofs, Domains, Ethics of Modelling

This chapter distinguishes logical consequence from physical consequence. We clarify how to report proofs alongside assumptions, how to police idealizations, and how to keep equations accountable to devices. The ethic is simple: say exactly what the proof proves, for which objects, under which limits, and how those objects meet the physical gates and the DoV.

16.1 What Proofs Can/Can't Do

Three things proofs *can* do for physics.

1. **Internal soundness.** Certify that a claimed consequence follows from stated axioms and definitions (e.g., stability of a generator, positivity of a kernel, uniqueness on a functional domain).

2. **Robust structure.** Identify invariants, monotones, and bounds that survive coarse-graining and therefore have a chance to be device-facing (e.g., Bochner positivity $\Rightarrow$ non-negative PSD).

3. **Counterexample economy.** Kill bad models early by proving that a desideratum is impossible under the chosen assumptions (saving experimental effort).

Three things proofs *cannot* do by themselves.

1. **Enlarge the DoV.** A theorem for an infinite, noiseless, continuum object does not certify behavior on finite, noisy, discretized devices unless a stability-at-finite link is proved.

2. **Turn abstractions into realities.** A limit object (plane waves, delta sources, exact symmetries) remains an idealization until a finite tolerance audit shows pointer-level stability.

3. **Bypass calibration and causality.** No algebraic argument replaces the need to declare units, to pass causality/no-signalling checks, and to provide a measurement bridge.

Principle 16.1: Proofs certify *logical*, not *physical*, consequence

A proof licenses the step "if the assumptions hold, then the statement holds" inside a formal structure. It does not license "the statement is true of nature" unless: (i) the objects are mapped to finite, unit-bearing, gauge-stable descriptors; (ii) the DoV is declared; and (iii) a model→pointer bridge is supplied.

Criterion 16.1: Admissible proof reporting (physics-facing)

When stating a theorem that supports a physical claim, include on the spot:

(a) **Object map.** What physical descriptors the abstract objects represent (symbols and units).

(b) **Scope.** The window of scales/approximations under which the hypotheses are expected to hold (DoV).

(c) **Stability.** A finite-tolerance statement or a reference to a stability-at-finite audit that links the limit to device scales.

(d) **Gates.** Which of S/G/TQ/D the theorem helps certify and which remain empirical or model-dependent.

title=Proof vs model vs measurement (roles)

Proof ensures no contradiction *within* the formal frame. **Model** ties symbols to descriptors with units and a DoV. **Measurement** tests the bridge and the gates on a device and a timescale. None of the three substitutes for another.

Worked contrast (limit vs finite). A proof that a continuum signal has non-negative PSD by Bochner's theorem is valuable. It becomes physics only after showing that the windowed, finite-length, discretized estimator $\widehat{S}(\omega)$ remains non-negative within uncertainty after deconvolution of the device window. The additional finite statement is what carries the proof to the pointer.

Principle 16.2: No idolatry of exactness

Exact equalities in the limit must be paired with inequalities at finite resolution. Reporting only the limit breeds reification; reporting the finite tolerance earns physical content.

16.1.1 Beyond Measurability

Certain mathematically meaningful statements outrun what current or even ideal devices can access. We neither outlaw such statements nor promote them to ontology. We bracket them and offer ways to bring them back inside physics.

Classes that tend to outrun measurability.

- **Global exact structure.** Properties that require access to infinite domains or exact continua (e.g., global sections on non-compact manifolds, exact symmetries across all scales).
- **Non-observables.** Gauge- or coordinate-dependent quantities with no invariant pointer meaning.
- **Undecidable/ill-posed claims.** Assertions that depend on limits whose order matters but is left unspecified (interchange of limit and integral, thermodynamic vs continuum limits).

Rule 16.1: How to domesticate the beyond-measurable

R1. Replace equalities by bounds. Trade "equals in the limit" for "lies in an interval at finite resolution" with a stated tolerance.

R2. Promote invariants. Re-express in gauge/coordinate-invariant terms or explicitly flag dependence as descriptive only.

R3. Fix limit orders. State the order of limits and prove stability under the chosen order on the declared DoV.

R4. Exhibit proxies. Provide operational proxies (counts, slopes, shares) that the device can stably estimate and that the theorem constrains.

Example (existence vs construction). A proof that a minimizer exists for an energy functional is not yet a device claim. A constructive scheme with error bounds that converge within the device window turns the existence result into a calibrated prediction.

Criterion 16.2: Green/Amber/Red for beyond-measurable statements

Green: rephrased as bounds on pointer proxies with units and a DoV. **Amber:** invariant, finite descriptors named; bridge sketched but tolerances pending. **Red:** gauge-dependent, limit-only, or order-of-limits–sensitive with no finite audit.

16.2 Idealization without Idolatry

Idealization is not the enemy of physics; it is how we think. But it becomes a liability the moment we *reify* it—treating a calculational device as if it were a physical subject. This section states working rules for using ideals (infinite limits, delta sources, exact symmetries, white noise, rigid backgrounds, linearization) without letting them outrun measurability, causality, or the stated domain of validity (DoV).

Principle 16.3: No reification of ideals

Use idealizations as tools to compute and to see structure; do not promote them to ontology. Each ideal must be accompanied by: (i) units and a calibration bridge; (ii) a DoV; and (iii) a stability-at-finite audit that shows the reported number converges within the device tolerance.

Criterion 16.3: Stability-at-finite (operational form)

Let Q be a quantity obtained via an ideal limit (infinite extent, continuum, $T \to \infty$, etc.). Provide a sequence of *finite* approximants $Q_{(h,N,T,\dots)}$ and a tolerance ε set by calibration. Report a window where

$$\frac{|Q_{(\cdot)} - Q'_{(\cdot)}|}{\max\{1, |Q_{(\cdot)}|\}} < \varepsilon \quad \text{whenever the finite cutoffs } (h, N, T, \dots) \text{ are tightened inside the DoV.}$$

If no such window exists, the idealization is *not* licensed for the claim.

Rule 16.2: Operational rules for safe idealization

I1. Declare the ideal early. Name where it enters (equation, lemma, fit) and why it is used (algebraic simplification, asymptotic picture, etc.).

I2. Pair every ideal with a finite audit. State the knobs (window length, grid step, packet width, cutoff Λ) and the device tolerance used to verify stability.

I3. Prefer invariants. Formulate results in gauge/coordinate invariants so the ideal does not hide chart artefacts.

I4. Publish the failure mode. Say explicitly what would break the ideal (e.g., nonlinearity, dispersion, memory), and how this would show up at the pointer.

I5. Record the abstraction budget. Count your ideals; each has a test plan and a date to be stress-tested or retired.

Typical ideals and how to domesticate them. *Plane wave* $\Rightarrow$ demonstrate that packet averages converge within tolerance on the reported band and time. *Delta source*

⇒ show that a compact source of radius a yields the same pointer within error for a below resolution. *Exact symmetry* ⇒ quantify the symmetry breaking and show that the reported number is stable under perturbations at the measured level. *White noise* ⇒ document bandwidth and verify that colored corrections do not change the reported statistic beyond uncertainty. *Rigid background* ⇒ demonstrate that the effective transport cone inferred from data coincides (within tolerance) with the background used.

title=Anti-patterns to avoid

"Taking limits first, declaring DoV later"; "reporting exact equalities without finite tolerances"; "using coordinate-dependent statements as if physical"; "fitting non-CP maps because they match a plot"; "negative PSD after windowing". These are red flags and should be rewritten or quarantined.

Worked contrast (white vs colored noise). If a response model assumes white noise to obtain a closed form, report the bandwidth and a colored-noise sensitivity check: vary the correlation time within calibration bounds and show the pointer statistic (slope, variance, share) moves by less than the quoted tolerance. If not, the claim is outside the model's DoV.

16.2.1 Reporting Assumptions

Assumptions are where most reification hides. We therefore require an explicit, machine-readable report for the assumptions behind each substantive claim or figure. Keep it short, but complete.

Rule 16.3: Assumption report (minimal fields per claim)

For each assumption A_i used in the derivation or fit, report:

A1. Name & role. What is assumed, and where it enters (equation, step, fit).

A2. Scale & units. Numbers with units; resolution/band where it is intended to hold.

A3. Gate touchpoints. Which of S/G/TQ/D the assumption affects (e.g., S via positivity, G via cone, TQ via CP, D via rate signs).

A4. Finite audit. The concrete knob(s) varied (packet width, cutoff, acquisition time) and the stability criterion used (tolerance).

A5. Failure signal. What pointer pattern would indicate that A_i is false on the declared DoV.

Table 16.1: Example assumption ledger (to be adapted per chapter).

Name/role	**Where it enters**	**Gates**	**Finite audit (knobs / tolerance)**	**Failure signal (pointer)**
Plane-wave approx.	Linear response; closed form for $S(\omega)$	S, G	Packet width; window length; $\varepsilon = 0.02$ slope tolerance under packet narrowing.	Slope change $> \varepsilon$ or bias under narrowing.
White-noise bath	Simplifies GKSL rates	TQ, D	Band limits; verify colored-noise correction does not flip signs.	Negative PSD band or EP sign flip.
Exact symmetry	Conserved quantity in derivation	S	Add perturbation δ; check invariant stability; drift $< O(\delta)$.	Observable drifts linearly with δ.
Rigid background	Transport cone fixed a priori	G	Compare inferred cone to background; misfit within tolerance.	Pre-front pointer change (no-signalling violated).

Criterion 16.4: Acceptance of an assumption set (per claim)

An assumption set is acceptable when: (i) all items have units, gates, and an audit; (ii) every idealization reports a stability-at-finite check; (iii) at least one falsifier is stated at the pointer level; (iv) the combined DoV is declared (intersection of individual DoVs). Otherwise the claim is either *borderline* (pending tests) or *outside* physics.

Author practice. Attach the assumption ledger as an appendix snippet to each chapter or major figure. Keep rows terse (one line each). Reviewers can request a script/notebook that reproduces the finite-audit plots. This keeps ideals as productive tools, not as accidental ontologies.

16.3 Style Guide for Equations

Equations are the sharpest interface between models and devices. This guide makes each displayed equation *accountable*: every symbol has units, the valid regime is declared, gauge status is clear, limiting ideals are audited at finite, and the relevant gates (S/G/TQ/D) are stated alongside the bridge to the pointer.

16.3.1 The Five Rules

Rule 16.4: The Five Rules (equation hygiene)

E1. Units & dimensions first. State units for all quantities in or under the equation (SI or explicitly named); include at least one dimensional check in the surrounding text.

E2. Declare the domain of validity (DoV). Right after the equation, declare scales/approximations/cutoffs and the triggers for end-of-validity (nonlinearity,

bandwidth edges, memory effects).

E3. Gauge status & invariants. Formulate claims in invariant quantities; if a gauge/coordinate choice is used in the derivation, flag it as descriptive and show the invariant final form or its gauge-insensitivity at the pointer.

E4. Limits only with finite audits. If the equation relies on an ideal limit ($T \to \infty$, continuum, plane waves, ∞), provide finite approximants and a stability-at-finite tolerance tied to calibration.

E5. Gate badges & bridge. For any equation used in predictions, state which gates it touches (S/G/TQ/D) and provide the model→pointer map (what the device reads, with calibration).

Badge line (author convention). After important equations, append a short "badge line" in prose with: *Units; DoV; Gauge; Limit audit; Gates; Bridge.* Keeping this text close to the equation prevents reification and drift.

title=Minimal badge line template

Units: $[X] = \cdots$. **DoV:** band/time/strain $\cdots$; end-of-validity: $\cdots$. **Gauge:** invariant form $\cdots$ (choice: $\cdots$ benign). **Limit audit:** windowed estimator with tolerance $\varepsilon = \cdots$. **Gates:** S (ok), G (cone g_{eff} ok), TQ (n/a or CP via GKSL), D (non-negative). **Bridge:** pointer channel(s) and calibration.

Worked example (linear response to pointer). Consider a calibrated linear channel with transfer function $H(\omega)$ driven by a stationary input ξ and additive detector noise n. The device reports the one–sided power spectral density (PSD)

$$S_{YY}(\omega) \;=\; |H(\omega)|^2\, S_{\xi\xi}(\omega) \;+\; S_{nn}(\omega)\,, \qquad \omega \in [\omega_1, \omega_2]. \tag{16.1}$$

Badge line. **Units:** $[S_{YY}] = (\text{unit})^2/\text{Hz}$, $[H] = \text{unit/unit}$. **DoV:** linear regime; band $[\omega_1, \omega_2]$; window length T with $\tau_B \ll T$. **Gauge:** invariant under reparametrization of state; reported in PSD (device–level invariant). **Limit audit:** estimate $\widehat{S}_{YY}$ via Welch/Multitaper; report stability vs T and binning with tolerance ε. **Gates:** S (PSD ≥ 0; Toeplitz positivity checked), G (no–signalling respected; front within g_{eff}), TQ (n/a if classical; CP needed if quantum reduction used), D (passivity $\Rightarrow$ non–negative absorption on the band). **Bridge:** analyzer calibration $\Pi : \{Y(t)\} \mapsto \widehat{S}_{YY}$ with window deconvolution.

Worked example (open quantum share vector). With a GKSL generator and POVM $\{E_r\}$, the time–averaged share in channel r is

$$s_r(T) \;=\; \frac{1}{T}\int_0^T \mathrm{Tr}\Big(E_r\,\rho(t)\Big)\,dt, \qquad \mathbf{s}(T) \xrightarrow[T\to\infty]{} \mathbf{p}. \tag{16.2}$$

Badge line. **Units:** shares dimensionless; $[E_r] = \mathbb{I}$. **DoV:** weak coupling; Markovian on T; detector linearity. **Gauge:** invariant under representation change; POVM fixed by calibration. **Limit audit:** Allan deviation of $\mathbf{s}(T)$ scales $\sim T^{-1/2}$ on the window. **Gates:** S (observables self–adjoint), TQ (CP via GKSL), D (entropy production ≥ 0), G (no–signalling under g_{eff} cone). **Bridge:** counter calibration $\Pi : \rho(t) \mapsto \{j_r(t)\} \mapsto \mathbf{s}(T)$.

Table 16.2: Common equation anti–patterns and how to fix them.

Anti–pattern	Fix
Equation with dimensionless symbols and no unit declarations	Add a symbol table with $[\cdot]$, run a dimensional check, and state calibration constants.
Gauge–dependent intermediate expression reported as "the effect"	Rewrite in invariants; demonstrate that pointer statistic is gauge–insensitive on the DoV.
Use of a limit (plane wave, $T \to \infty$) with no finite audit	Provide finite approximants and a stability–at–finite tolerance tied to device resolution.
Non–CP effective map fitted because it "matches the plot"	Constrain the fit to the CPTP cone or enlarge the model; otherwise quarantine outside TQ.
Negative PSD after windowing	Recalibrate windows, extend acquisition, enforce kernel positivity (S gate).

Anti–patterns (and fixes).

Author practice (checklist per displayed equation).

- Symbol/units table near first use; state constants and their units.
- DoV sentence: scales, approximations, cutoffs, and end-of-validity triggers.
- Gauge status: invariant form or explicit statement that a choice is descriptive only.
- Limit audit: which finite knobs were tightened; reported tolerance ε; stability plot referenced.
- Gates touched and evidence: S/G/TQ/D in one line.
- Bridge: which pointer channel reads the quantity and how it is calibrated.

Part IX

Objections and Replies

A disciplined philosophy of physics should welcome objections: they pressure–test our commitments and sharpen the workflow. This part addresses three familiar challenges. First, that our device–anchored stance collapses into mere instrumentalism. Second, that only *structure* (not things) is real. Third, that probability in our account simply re–assumes Born. In each case we reply by showing how the S/G/TQ/D gates, domains of validity (DoV), and the model→pointer bridge distinguish empty rhetoric from admissible physics.

Chapter 17

Instrumentalism, Structural Realism, Born

This chapter treats the first objection in detail (§17.1), and sets up the responses to structural realism and Born in the next sections. Our replies are not slogans; they are *operational*: they specify what an author must show so that a claim earns physical content without reifying scaffolding.

17.1 Just Instrumentalism?

The objection: if we only accept what reaches a pointer with units, aren't we just "instrumentalists" dressing up curve–fitting with grand language?

Short answer. No. Our stance is *gauge realism with device anchoring.* It is realist about what survives redundancy (invariance) and stabilizes under coarse–graining on a declared DoV, and it is strict about how symbols earn that status (bridge, gates, stability–at–finite). This is stronger than instrumentalism (which is content with prediction alone) and stricter than naive realism (which reifies descriptions).

Principle 17.1: Anchored realism

What is *physically* real in a model is what (i) is expressed in invariant, finite descriptors; (ii) persists under admissible reparametrizations and coarse–grainings; and (iii) is tied by a calibration map to pointer statistics on a stated DoV. Prediction is necessary but not sufficient; invariance and stability provide the ontic anchor.

Criterion 17.1: Not–just–instrumental test

A claim avoids mere instrumentalism when, besides predictive accuracy, it supplies: (a) an invariant formulation; (b) a stability–at–finite audit linking any limit ideal to device scales; and (c) a gate status line (S/G/TQ/D) tied to the DoV. Lacking any of (a)–(c), the claim is at best instrumental.

Why invariance matters. Two algebraically different descriptions can make identical predictions within noise. Instrumentalism stops there. Anchored realism asks which parts of those descriptions survive re–gauging and coarse–graining; only *those* parts are licensed as candidates for ontology. This is how we keep "coordinate artefacts" and "exact limits" from sneaking into metaphysics.

Worked contrast (schematic). A foliation–dependent stress "density" is used to argue for a physical effect. The instrumentalist says: it fits the plot. Our audit: rewrite in curvature or other invariants; show the effect persists and passes S/G/TQ/D on the DoV; show stability under packet–vs–plane–wave replacement. If not, the "effect" was an artefact. If yes, we have *more than* fit: we have an anchored, realist content.

title=What distinguishes our stance from two extremes

Vs naive realism: we refuse to reify coordinates/limits; ontology is earned by invariance and stability.
Vs instrumentalism: we demand *explanation* as pattern under symmetry/coarse–graining, not mere fit; we require gates and a DoV, which constrain future predictions beyond curve–fit.

17.1.1 Gauge Realism

Gauge realism is the backbone of anchored realism: the world does not change when we change our bookkeeping.

Definition 17.1: Gauge realism (operational)

A report is *gauge realist* when its physically loaded quantities are expressed as invariants under the admissible redundancy group (gauge, coordinates, representation), and when any intermediate choice (gauge fixing, chart, picture) is flagged as scaffolding that disappears at the pointer.

Three tests of gauge realism.

(a) **Invariant restatement.** Provide an explicitly invariant form of the claim (e.g., scalars, eigenvalues, spectra, equivalence classes).

(b) **Pointer insensitivity.** Show that pointer statistics are unchanged (within tolerance) under allowed reparametrizations on the declared DoV.

(c) **Cone consistency.** For propagation claims, the effective cone inferred from data is independent (within uncertainty) of descriptive choices.

Rule 17.1: Gauge–realist reporting (author checklist)

G1. State the redundancy group and the invariant quantities used.

G2. If a gauge/coordinate is fixed for algebra, flag it; give the invariant final expression and the pointer–level check.

G3. Record how re–gauging affects intermediate symbols and why the bridge map to the pointer is insensitive on the DoV.

Mini examples (one–liners). *GR:* horizon vs singularity—report curvature invariants, not chart coefficients; pass S/G, neutral on TQ/D. *QFT:* renormalized masses/couplings—report scheme–independent observables (scattering lengths, spectral residues) and stability across scales on the DoV. *Open systems:* present GKSL rates tied to calibrated baths; ensure that a change of representation leaves the POVM–based pointer shares invariant.

Criterion 17.2: When "structure only" is enough, and when it is not

Structure (relations, symmetries) is sufficient *if* it determines invariant, finite pointer predictions on the DoV. Otherwise, structure without a bridge is mathematics, not physics, and must be bracketed as motivational.

Bottom line. We are not "just" instrumentalists. We are realists about invariant, stable patterns that reach the device through a disciplined bridge and survive the gates. The rest of this chapter applies the same discipline to structural realism and to the status of Born in our pipeline account.

17.2 Only Structure Is Real?

Structural realism claims that what survives theory change is not "things" but *structure*: relations, symmetries, and morphisms. Our stance is close but stricter: *structure earns reality only when it anchors to invariant, finite, device–facing predictions on a declared DoV*. Unanchored structure is excellent mathematics and may guide modelling; it is not yet physics.

Principle 17.2: Anchored structural realism

A structural claim is *physical* iff it is expressed in gauge/coordinate invariants, carries units (or dimensionless ratios tied to units elsewhere), passes the S/G/TQ/D gates on its domain of validity (DoV), and furnishes a model→pointer bridge. Otherwise it is structural *scaffolding*.

Why structure alone is not enough. A category of models with pretty morphisms, or a symmetry group acting on equations, may organize our thinking without yet fixing what a device reads. Anchoring requires (i) invariants that survive reparametrization, (ii) stability under coarse–graining, and (iii) calibration to pointers. Only then does "structure" constrain reality rather than describe our bookkeeping.

Criterion 17.3: Adequacy of a structural claim

A structural claim is adequate for physics when it supplies:

(a) **Invariants with units.** Concrete invariants (e.g., spectra, indices, curvature scalars, topological numbers) tied to dimensional descriptors or to calibrated ratios;

(b) **A stability report.** The invariant is stable at finite resolution under coarse–graining and noise on the stated DoV;

(c) **A gate line.** S/G/TQ/D checks relevant to how the structure is used (e.g., positivity for spectral structure, cone/no–signalling for causal structure);

(d) **A bridge.** A calibrated map from the invariant to pointer statistics (plots, counts, shares) with an uncertainty budget.

title=Structure: when to bracket, when to promote

Bracket (mathematics, not physics): global exact symmetries on continua with no finite audit; gauge–dependent "structures"; order–of–limits–sensitive claims without fixed order. Promote (physics): invariant quantities whose pointer proxies are stable on the DoV and that pass the gates.

17.2.1 Anchored Patterns

We now show how to turn bare structure into *anchored patterns*: stable, invariant summaries with units that reach a pointer. The recipe is uniform across domains.

Rule 17.2: Recipe: from structure to anchored pattern

P1. Name the structure. Symmetry G, topology class, RG fixed point, equivalence relation, cone field, spectral decomposition, etc.

P2. Extract invariants. Identify invariants that the structure fixes: indices, eigenvalue sets, conserved currents, exponents, cones, residues.

P3. Attach units & DoV. Express invariants as unit–bearing descriptors or ratios tied to calibrated scales; declare the window (DoV).

P4. Supply the bridge. Give a model→pointer map and the statistic to be reported (slope, knee, share, count, PSD level) with calibration.

P5. Pass the gates. State S/G/TQ/D checks relevant to the pattern (e.g., positivity for spectral patterns; no–signalling for cone patterns; CP for open quantum patterns; non–negative EP for irreversible patterns).

P6. Falsify explicitly. Provide a pointer–level signature that would disconfirm the pattern on the DoV.

Mini–atlas of anchored patterns (schematic).

- **GR (cone structure).** Structure: Lorentzian cone field. Invariant: light–cone order; curvature scalars. Bridge: pump–probe fronts; arrival–time isochrons. Gates: G (cone/no–signalling), S (positive PSD for responses). Falsifier: pre–front pointer change in the no–signalling test.

- **QFT (RG fixed point).** Structure: scale invariance at a fixed point. Invariant: critical exponents $(\nu, \eta, \dots)$. Bridge: log–log slopes in two–point functions; finite–size scaling of spectra. Gates: S (positive spectral density), D (coarse–grained kinetics non–negative). Falsifier: slopes drift beyond tolerance under binning/size changes.

- **Topological phase (band structure).** Structure: Chern number C. Invariant: integer C. Bridge: quantized Hall conductance plateaus; edge–mode counts. Gates: S (no negative–energy artefacts), G (causal transport), TQ (CP for open probes). Falsifier: plateau breaks within DoV without disorder explanation.

- **Open quantum pattern.** Structure: GKSL semigroup class. Invariant: steady–state manifold dimension; Lindblad spectrum gaps. Bridge: energy–share vector $\mathbf{p}$; gap shows up as mixing time. Gates: TQ (CP), D (entropy production), S (observable spectra). Falsifier: non–CP tomography or negative EP on the window.

Worked example (RG exponent as anchored pattern). Structure gives $G(r) \sim r^{-(d-2+\eta)}$ at criticality. *Anchoring:* report the slope $\widehat{\eta}$ from finite windows $r \in [r_1, r_2]$ with stability vs binning; units via lattice spacing a; DoV via $r_1 \gg a$ and $r_2 \ll L$; S gate via Bochner positivity of G's PSD; D gate via non–negative coarse–grained entropy in the kinetic model used for data collapse. Falsifier: slope drift $> \varepsilon$ under window shift with calibration held fixed.

Criterion 17.4: Acceptance test for an anchored pattern

Accepted iff: (i) invariant defined and unit–bearing; (ii) stability–at–finite documented; (iii) gate line given; (iv) falsifier named. Otherwise the pattern remains motivational.

title=Common misfires (and quick fixes)

Misfire: quoting a symmetry to explain a plot without an invariant. **Fix:** specify the invariant (exponent/index) and report the finite–window estimate. **Misfire:** "causal structure" shown with group–velocity curves. **Fix:** switch to front–velocity and no–signalling interventions. **Misfire:** topological claim with no unit–tied observable. **Fix:** give the quantized response and its calibration.

Bottom line. We welcome structure—not as metaphysics, but as a generator of *anchored patterns* that survive invariance tests, stabilize at finite resolution, pass the gates, and land on calibrated numbers. Only then does "structure" count as real for physics.

17.3 Is Born Assumed?

A standing worry is that any account which ends up with $p(r) = \mathrm{Tr}(E_r\,\rho)$ has merely *assumed* Born in disguise. Our answer is operational: we assume only (i) calibrated instruments that act as completely positive (CP) maps on the declared domain of validity (DoV), (ii) open dynamics that are CP and causal on the reporting window, and (iii) finite, unit–bearing pointer summaries. From these, the Born form *emerges* as the *stable long–time share* of outcomes in a measurement pipeline, not as a postulate about probabilities.

Principle 17.3: Born as a stability statement, not a postulate

For a calibrated instrument with POVM elements $\{E_r\}$ and CP instrument maps $\{\mathcal{I}_r\}$ acting on a system with open dynamics that is CP on the DoV, the long–time outcome shares converge to

$$p(r) \;=\; \mathrm{Tr}(E_r\,\rho_{\mathrm{eff}}),$$

where ρ_{eff} is the effective pre–readout state delivered by the pipeline. The statement is about *frequencies as limits of pointer counts* under declared resolution and mixing, not an a priori probabilistic axiom.

title=What we *do* and *do not* assume

We assume: CP/CPTP structure (Section 14.3), calibrated pointers, a DoV, and causal openness (Sections 14.2, 14.4). We do *not* assume: a primitive rule equating squared amplitudes with probabilities, collapse postulates, or ideal detectors outside calibration.

17.3.1 Pipeline Derivation

We now sketch an end–to–end derivation that lands on the Born form as a *frequency claim*.

Step 1: Instrument from unitary probe and readout. Let S (system) couple to a meter M via a calibrated unitary U during a shot of duration Δ, then a classical backend coarse–grains the meter into outcomes $r \in \mathcal{R}$. The completely positive *instrument* maps are

$$\mathcal{I}_r(\rho) \;=\; \sum_k M_{r,k}\,\rho\,M_{r,k}^\dagger, \qquad E_r \;=\; \sum_k M_{r,k}^\dagger M_{r,k}, \qquad \sum_r E_r = \mathbb{I}, \tag{17.1}$$

so the total channel is $\Phi(\rho) = \sum_r \mathcal{I}_r(\rho)$ and E_r are the POVM elements fixed by calibration.

Step 2: Open evolution between shots. Between shots we allow CP, time–local open dynamics

$$\dot{\rho}_t = \mathcal{L}_t(\rho_t) \qquad \text{(GKSL on the DoV)}, \tag{17.2}$$

or its discrete propagator $\mathcal{E}_\Delta = \exp\int_t^{t+\Delta}\mathcal{L}_\tau d\tau$.

Step 3: The one–cycle channel and mixing. One acquisition cycle is

$$\mathcal{T} \;=\; \Big(\sum_{r\in\mathcal{R}} \mathcal{I}_r\Big) \circ \mathcal{E}_\Delta, \tag{17.3}$$

which is CPTP. On the reported window we assume *mixing*: $\mathcal{T}$ has a unique stationary state ρ_* with spectral gap > 0 on the DoV, so $\rho_{n+1} = \mathcal{T}(\rho_n) \to \rho_*$ as $n \to \infty$ regardless of ρ_0.

Step 4: Outcome process and long–time shares. Let $X_n^{(r)} \in \{0,1\}$ indicate that outcome r occurred on cycle n. Conditioned on ρ_n and the instrument calibration,

$$\mathbb{E}[\, X_n^{(r)} \,|\, \rho_n \,] \;=\; \mathrm{Tr}(E_r\, \rho_n). \tag{17.4}$$

Define the empirical share $s_r(N) = \frac{1}{N}\sum_{n=1}^{N} X_n^{(r)}$. Then by mixing and a martingale/ergodic law of large numbers,

$$s_r(N) \;\xrightarrow[N\to\infty]{\text{a.s.}}\; \mathrm{Tr}(E_r\, \rho_*) \quad \text{and} \quad \sqrt{N}\Big(s_r(N) - \mathrm{Tr}(E_r\, \rho_*)\Big) \;\Rightarrow\; \mathcal{N}(0, \sigma_r^2), \tag{17.5}$$

with σ_r^2 set by the instrument and the gap (calibration reports it).

Step 5: The reported probabilities. The object we report is thus

$$p(r) \;:=\; \lim_{N\to\infty} s_r(N) \;=\; \mathrm{Tr}(E_r\, \rho_*)\,, \tag{17.6}$$

which is the Born form with *effective* pre–readout state ρ_* delivered by the pipeline. If the acquisition is fast or the environment weak, ρ_* is practically the prepared ρ; otherwise it is the mixed state after $\mathcal{E}_\Delta$ and any pre–processing in $\mathcal{T}$.

Comments and variants. (i) *Projective/QND limit:* if $M_{r,k}$ are projectors and $\mathcal{E}_\Delta$ vanishes, (17.6) reduces to textbook Born. (ii) *Inefficiency and dead time:* observed $\tilde{E}_r = \eta_r E_r + B_r$ with known efficiencies $\eta_r \in (0,1]$ and background $B_r \geq 0$; calibration inverts this to recover E_r (within uncertainty). (iii) *Correlated shots:* if $\mathcal{T}$ is not mixing on the window, report $s_r(N)$ with its Allan deviation and restrict claims (outside our Born stability principle).

Criterion 17.5: When you may claim a pipeline derivation of Born

You may assert (17.6) only if you provide:

B1. Calibration: instrument maps $\{\mathcal{I}_r\}$ or POVM $\{E_r\}$ with units and uncertainties;

B2. Openness: a CP description of $\mathcal{E}_\Delta$ (or a declared limit where it is negligible);

B3. Mixing evidence: spectral gap or empirical stabilization of $s_r(N)$ with $V(N) \sim N^{-1/2}$ on the DoV;

B4. Gate line: S (positivity of E_r), TQ (CP of $\mathcal{T}$), D (non–negative EP on the window), G (no–signalling in interventions).

Lacking any item, downgrade the statement to a calibration–dependent fit, not a derivation.

title=How this rebuts "Born is assumed"

We never posit $p(r) = \|\Pi_r \psi\|^2$. We posit CP instruments and openness on a DoV; we then *prove* that long–time pointer shares stabilize to $\mathrm{Tr}(E_r \rho_{\mathrm{eff}})$, with ρ_{eff} and E_r both calibrated, and we exhibit the $N^{-1/2}$ approach in the data. The Born form is the *name* for that stabilized share, not a postulate.

Typical failure modes (and what they mean). Non–CP fitting (negative POVM elements), strong non–Markovian memory that breaks mixing on the window, detector saturation that corrupts $X_n^{(r)}$, or causal cross–talk that violates no–signalling: each breaks an antecedent of the pipeline and invalidates the Born stability claim for that DoV. The fix is to restore CP/causality or restrict the scope explicitly.

On this reading, "Born" is a compact label for a robust frequency statement about calibrated devices operating on an open system under well–stated conditions. It is *derived* from the pipeline assumptions you already owe to the reader, not an axiom added on top.

Chapter 18

What About Gauge?

Gauge freedom is bookkeeping freedom: it is how we label states and fields so that calculations are tractable. Reality, in our stance, sits in what survives all admissible relabellings—*invariants* and, when appropriate, entire *equivalence classes* of descriptions. This chapter separates what is real from what is descriptive, provides tests and reporting rules to avoid reification, and ties the discussion to device–facing pointers and the gates.

18.1 Reality vs Description

We begin by making the distinction operational across the main arenas where gauge freedom appears: internal gauge theories, diffeomorphism invariance, and effective descriptions (e.g., emergent cones g_{eff}).

Definition 18.1: Gauge transform vs physical state

A *gauge transformation* is a relabelling action G on the model's descriptive variables that leaves all admissible pointer predictions unchanged. Two descriptions related by $g \in G$ represent the *same* physical state.

Internal gauge (example: electromagnetism). Potentials A_μ transform as $A_\mu \mapsto A_\mu + \partial_\mu \lambda$. Invariants include the field strength $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$, Lorentz scalars $F_{\mu\nu}F^{\mu\nu}$ and $\tilde{F}_{\mu\nu}F^{\mu\nu}$, and Wilson loops $W(C) = \operatorname{Tr} \mathcal{P} \exp\Big(i \oint_C A_\mu dx^\mu\Big)$. Device bridges: interference phase shifts (Aharonov–Bohm), spectral lines, force readouts—each tied to invariants.

Geometry (example: GR). A diffeomorphism $x^\mu \mapsto x'^\mu(x)$ changes metric components $g_{\mu\nu}$ but preserves curvature invariants (e.g., R, $R_{\mu\nu}R^{\mu\nu}$, $R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$) and causal order. Device bridges: time–delay, lensing angles, redshift—quantities definable without chart commitment.

Effective cones and media. In media we infer an effective cone or metric–like descriptor $g_{\text{eff}}[\mu\nu]$ from fronts and dispersion. Chart choices or parameter redefinitions are descriptive; the *isochron pattern* and front–velocity bounds are the invariants tied to pointers.

Principle 18.1: No ontology for gauge choices

Coordinate charts, gauge fixings, phase conventions, and unphysical polarizations are scaffolding. Only gauge–invariant quantities or equivalence classes may carry ontic weight in our claims; everything else is to be explicitly marked as *descriptive*.

Criterion 18.1: Gauge–invariant reporting (physics–facing)

Any claim with gauge freedom must include:

(a) **Invariants.** An explicitly invariant statement (scalar, spectrum, Wilson loop, causal order, equivalence class).

(b) **Bridge.** A calibration from the invariant to pointer statistics.

(c) **DoV.** The scales/approximations where the invariant controls the pointer.

(d) **Gate line.** S/G/TQ/D checks relevant to how the invariant is used.

title=Reality vs description—quick map

Real (candidates): spectra of self–adjoint observables; curvature and field invariants; Wilson loops; scattering phases; cone order and fronts; share vectors stabilized at long time.
Descriptive (scaffolding): gauge potentials A_μ without loops; metric components in a chart; lapse/shift as coordinates; gauge parameters; unphysical polarizations; renormalization scheme parameters (absent observables).

Device test for gauge insensitivity. Given a predicted pointer summary P and a family of gauge choices $g \in G$, fit the same data (or simulate) under multiple g and report

$$\Delta_G(P) \;=\; \sup_{g,g' \in G} \frac{|P^{(g)} - P^{(g')}|}{\max\{1, |P|\}},$$

with the same calibration. Passing threshold $\Delta_G(P) < \varepsilon$ (set by device tolerance) is part of the acceptance metric.

Rule 18.1: Author checklist (Reality vs Description)

R1. State the redundancy group G and the allowed gauge/coordinate moves.

R2. Present the claim in invariants; if a gauge is fixed for algebra, flag it and show gauge–insensitive pointer predictions.

R3. Declare the DoV and list gate checks (S for positivity, G for cone/no–signalling, TQ for CP if reduction is used, D for entropy sign).

R4. Provide a $\Delta_G(P)$ report or an analytic argument for insensitivity.

18.1.1 Anti-Reification

Reification is the slide from "useful description" to "existent thing". Gauge contexts are especially prone to it. Here are the main traps and cures.

Traps (with fixes).

- **Potentials as causes.** Treating A_μ as causal agent independent of loops. *Fix:* report Wilson loops or field invariants; calibrate via interference phase.
- **Chart artefacts in GR.** Reading off "singular behavior" from metric components at horizons. *Fix:* use curvature invariants and causal structure; see the horizon vs singularity audit.
- **Gauge modes as signals.** Reporting longitudinal/ghost modes as detected fields. *Fix:* enforce spectral safety (no negative/ghost modes) and use gauge–fixed algebra only as scaffolding.
- **Scheme dependence as physics.** Taking renormalization scheme parameters as observables. *Fix:* report scheme–independent combinations (scattering lengths, residues) and provide stability across scales.
- **Effective metric overreach.** Promoting a fitted g_{eff} parameter to "geometry" beyond the band/DoV. *Fix:* tie claims to front and isochron data; declare DoV explicitly.

Principle 18.2: Anti–reification workflow

Before attaching ontic language to a symbol: (i) rewrite the claim in invariants; (ii) verify gauge/coordinate insensitivity at finite resolution; (iii) pass the gates on the declared DoV; (iv) name a pointer–level falsifier.

Rule 18.2: Minimal template: writing non–reifying claims

Claim (invariant, with units). .....
Descriptive choices used. (gauge/coordinate, purpose) .....
Gauge–insensitivity. $\Delta_G(P) < \varepsilon$ (evidence) .....
DoV & gates. scales, S/G/TQ/D status .....
Bridge & falsifier. pointer map; disconfirming signature .....

Worked micro–example (Aharonov–Bohm). *Structure:* $U(1)$ gauge with nonzero flux through a solenoid; A_μ is gauge–dependent. *Invariant:* loop phase $\Phi = \oint A_\mu dx^\mu$ (mod 2π) and interference shift. *Bridge:* calibrated fringe displacement vs injected flux. *DoV & gates:* linear optics, stable cone; S (PSD non–negative), G (no–signalling), TQ (n/a if classical illumination), D (passive optics). *Anti–reification:* no claim that "A_μ exists as a force"; the claim is about the invariant phase captured by the pointer.

Bottom line. Gauge freedom is a feature, not a bug. It lets us compute flexibly while keeping ontology disciplined: what is real are the invariant, finite, stable patterns that make it to the device and clear the gates; the rest is description we are free to choose—and obliged to label.

18.2 Pragmatics

This section turns the previous chapter's norms into day-to-day writing habits. It is a style-and-tools kit for authors and reviewers: how to phrase claims so they are gauge-real, device-facing, and gate-aware; how to caption figures; and how to avoid reification by construction.

18.2.1 Writing Non-Reifying Claims

A non-reifying claim is one that says exactly what the device can in principle read, in invariant terms, on a declared domain of validity (DoV), with the relevant gates (S/G/TQ/D) named. Use the following workflow to draft such claims, and to refactor legacy text.

Seven-step drafting workflow.

P1. Start with the pointer. What statistic will be reported (count, share, slope, PSD level, arrival-time front)? Name it and its units.

P2. Name the invariant. Express the claim in gauge/coordinate invariants (curvature scalars, spectra, Wilson loops, front velocity, etc.).

P3. Attach the bridge. One sentence mapping model descriptors to the pointer statistic (calibration reference, estimator).

P4. Declare the DoV. Scales, band, resolution, approximations, and what ends validity (e.g., nonlinearity, memory, bandwidth edges).

P5. Audit ideals at finite. If a limit was used (plane wave, continuum, $T \to \infty$), state the finite knobs and the stability tolerance.

P6. Gate line. State which gates are touched and the evidence: S (positivity), G (cone/no-signalling), TQ (CP), D (non-negative EP).

P7. Name a falsifier. A pointer-level pattern that would disconfirm the claim on the DoV.

Rule 18.3: Do/Do not (quick rewrite guide)

Do not say (reifying/descriptive)	**Say instead (non-reifying)**
"The potential causes the shift."	"The loop phase (Wilson loop) is $2\pi\Phi/\Phi_0$, shifting the calibrated fringe by Δx on the stated DoV."
"The metric component blows up at the horizon."	"Curvature invariants are finite at $r = 2M$; fronts respect the effective cone. The blow-up is chart-only."
"Information travels superluminally in the medium."	"Group velocity exceeds $c_\star$ on a narrow band, while the front velocity and no-signalling interventions remain within the cone."
"White noise bath holds."	"Within band $[\omega_1, \omega_2]$, colored corrections shift the pointer statistic by $< \varepsilon$ (finite audit)."

Caption checklist for figures. Every figure that underwrites a claim should include, in the caption or nearby: (i) units; (ii) DoV; (iii) gauge/invariant statement; (iv) finite-audit note; (v) gate line; (vi) bridge/estimator and calibration.

title=Minimal caption template (fill the brackets)

Units: $[Y] = \cdots$. **DoV:** band/time/scale $\cdots$. **Invariant:** $\cdots$ (gauge choices are descriptive). **Finite audit:** knobs (window T, grid h, packet width); tolerance

$\varepsilon = \cdots$. **Gates:** S $(\cdots)$, G $(\cdots)$, TQ $(\cdots)$, D $(\cdots)$. **Bridge:** Π : model $\to$ pointer; calibration ref $\cdots$.

Micro-refactors (three domains). *GR.* Replace chart statements with curvature scalars or causal order; report front maps and time delays with uncertainties. *QFT.* Replace scheme parameters by scattering lengths/spectral residues; state the energy window and RG stability. *Open quantum.* Report POVMs, GKSL rates, and share vectors; show CP of reduced maps and non-negative entropy production on the window.

Criterion 18.2: Pass/fail for a written claim (lint rules)

A sentence advertising a physical effect is *admissible* iff it contains, or directly references, all of: units; DoV; invariant form; finite-audit note if a limit is used; gate line; bridge. Missing items downgrade the sentence to a mathematical or descriptive remark.

Template: Non-reifying claim (one-liner + badge)

On band $[\omega_1, \omega_2]$ the invariant damping rate Γ of mode k_* satisfies $0 \leq \Gamma \leq \Gamma_{\max}$ (s^{-1}).
Badge. DoV: linear, $T = 10$ ms, strain $< 10^{-3}$. Invariant: PSD level from calibrated response. Finite audit: Welch T sweep, tol $\varepsilon = 0.02$. Gates: S (PSD ≥ 0), G (fronts within g_{eff}), TQ (CP via GKSL), D (EP ≥ 0). Bridge: analyzer calibration Π.

Common pitfalls (and fast cures).

- *Gauge words in the claim.* Cure: rewrite in invariants and move the gauge choice to methods.
- *Limit-only equalities.* Cure: add a finite audit with knobs and tolerance; otherwise bracket as mathematical.
- *Gate-free predictions.* Cure: add the gate line or restrict scope to a modelling lemma (not a physical claim).

Reviewer practice. When a manuscript mixes descriptive and physical statements, request a "gate card" (Section 15.2) per claim, or annotate with the lint rule **??**. Most reification disappears once authors supply invariants, DoV, finite audits, and gate lines.

The pay-off of this pragmatics is twofold: (i) claims become legible and checkable across subfields; (ii) abstraction keeps its power without being mistaken for ontology.

Part X

Program, Notes, Appendices

Part rationale. This part distills the book's discipline into short publishable notes and working templates. Each note is meant to be executable by a small team: state the descriptor, run the S/G/TQ/D gates, declare the domain of validity (DoV), and give a device–level bridge. The emphasis is on *method*, not on advocacy.

Chapter 19

A Research Program

The program is modular. Each module takes a notorious point of confusion (singularities, gauge "reality", probability, internal time), rewrites it in the book's operational language, and produces a minimal template that other groups can adapt. The goal is cumulative auditability rather than a single decisive theorem.

19.1 From Singularity to DoV

Aim. Turn any claim of "singularity" into (i) an invariant diagnosis, and (ii) a *domain-of-validity* statement with an *effective replacement* that preserves the S/G/TQ/D gates and offers a pointer bridge. The outcome is not a metaphysical "resolution" but a physically *admissible* model on a declared scope.

Principle 19.1: Singularity → DoV.

A divergence or blow-up is either (a) a removable *coordinate* artifact, or (b) an *end-of-validity* flag for the current descriptors/dynamics. In case (b), one must (i) state the DoV boundary with a stability-at-finite audit, and (ii) supply an *effective replacement* that keeps invariants finite and passes S/G/TQ/D on the target timescale.

Three recurring arenas.

- **GR (geometry).** Distinguish coordinate effects from invariant blow-ups (e.g. curvature scalars such as $K = R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$). Horizons are chart–dependent; a diverging K is an end-of-validity alarm.

- **QFT (UV).** Treat ultraviolet divergences as a *cutoff dependence* of intermediate quantities. Physical predictions must be *stable at finite* cutoff Λ on the reported accuracy:
$$\frac{|O(\Lambda+\delta)-O(\Lambda)|}{\max\{1,|O(\Lambda)|\}} < \varepsilon \quad \text{for} \quad \Lambda \geq \Lambda_0.$$

When not satisfied, change the effective theory or the observable.

- **Continuum PDE.** Shock formation or gradient blow-up in ideal models (e.g. inviscid limits) indicates missing dissipation/coarse-graining. Replace with a viscous/stochastic effective equation and report the *window* where its parameters were calibrated.

title=Minimal "Singularity → DoV" workflow (one page)

1. Invariants. Compute a chart/gauge invariant that diagnoses the issue (curvature scalars in GR; gauge-invariant spectra/poles in QFT; conserved densities in PDE).
2. Coordinate triage. Exhibit a chart where the effect disappears (if it is coordinate) or show an invariant blow-up (if it is DoV).
3. Stability-at-finite. Specify the finite approximants (cutoff, mesh, window) and demonstrate stability on the stated tolerance.
4. DoV boundary. State the boundary in *units* (scale, curvature, frequency, resolution).
5. Replacement. Provide an effective model (next subsection) that preserves S/G/TQ/D and a pointer bridge (how the device reads it).

Table 19.1: From "singularity" talk to operational statements.

Arena	Invariant cue	DoV boundary (units)	Effective replacement
GR (geometry)	$K\uparrow$	$K > K_\star$ (reported)	Core model with bounded stress; matched exterior; declare chart conditions.
QFT (UV)	$O(\Lambda)$ unstable	$\Lambda < \Lambda_\star$	EFT at scale $\Lambda_\star$ with running couplings; report window & tolerance.
PDE (shocks)	$\|\nabla u\|\uparrow$	$\Delta x < \Delta x_\star$	Viscous/stochastic closure; weak-solution regime; calibration of numerical dissipation.

19.1.1 Effective Replacements

Design constraints. Any replacement must:

C1 preserve *invariants* relevant to the claim (gauge/coordinate independence);

C2 pass the **S/G/TQ/D** gates on the stated timescale;

C3 include *units*, a *DoV declaration*, and a *pointer bridge*.

Four families of replacements.

1. **Coordinate surgery (GR).** Replace a bad chart with a regular one (Eddington–Finkelstein/Kruskal near horizons); keep the physical claim in curvature scalars or geodesic incompleteness tests.

2. **Wilsonian coarse-graining (QFT).** Introduce an explicit cutoff Λ and a scale–dependent action S_Λ; report *running* couplings and verify stability of observables at finite Λ within tolerance.

3. **Dissipative/stochastic closure (PDE/fluids).** Add the leading symmetry–allowed terms (viscosity, diffusion, noise) that render spectra finite and select weak solutions; calibrate coefficients from data.

4. **Open–system embedding (quantum/relativistic media).** Replace ill–posed closed dynamics by an effective generator that respects causality and complete positivity (e.g. a GKSL form consistent with the inferred cone g_{eff}).

Rule 19.1: "Replacement" author checklist

1. **State the invariant symptom** (what blew up, in what units).
2. **Declare the DoV boundary** as a number with uncertainty and calibration.
3. **Specify the replacement** (coordinate, EFT, closure, open-system) and why it is symmetry/gauge compatible.
4. **Gate line (S/G/TQ/D).** One sentence each: spectrum non-negativity; causal cone; complete positivity under reductions; non-negative entropy/linewidth.
5. **Pointer bridge.** Which device observable reads the replacement parameters?

Two sketched notes (ready to spin off).

- *From curvature blow-up to core model.* Use an equation-of-state ceiling to bound invariants ($K \leq K_\star$), match to Schwarzschild/Kerr exterior, and quantify the window where ringdown spectra are unchanged to within tolerance.

- *From UV divergence to calibrated EFT.* Choose a process (e.g. $2 \to 2$ scattering), define $O(\Lambda)$, show stability on $[\Lambda_0, 2\Lambda_0]$ for ε set by the device, and publish the gate line with a replication script.

19.2 Gauge-Invariant Existence

Aim. To say that something *exists* physically we require more than a name in a model. Existence is earned when a descriptor (i) is invariant under admissible gauge/coordinate

changes, (ii) survives coarse–graining (stability at finite), and (iii) admits a device–level bridge. This section turns that slogan into a workflow and applies it in QFT and GR.

Criterion 19.1: Gauge-invariant existence.

A claimed entity/process is accorded physical existence in this volume if its descriptor D satisfies:

1. **Invariance:** D is unchanged under admissible gauge/coordinate transformations (or depends only on invariants such as spectra, curvature scalars, holonomies).
2. **Stability-at-finite:** reported values of D are stable under the finite approximations inherent to the experiment or computation (cutoffs, bins, windows).
3. **Bridge:** there is a clear map from D to pointer observables (calibration curve, counting rule, or frequency extraction).
4. **Gates:** the claim passes S/G/TQ/D on the stated timescale.

title=Existence workflow (author checklist)

1. Name the invariant. Give D as a gauge/coordinate invariant (e.g. a pole location, a curvature scalar, a holonomy).
2. Declare the DoV. Units, resolution, and stability-at-finite audit.
3. Provide the bridge. How the device extracts D (fit protocol, spectral line, geometric observable).
4. Gate line. One sentence each for S/G/TQ/D.
5. Non-examples. State nearby gauge-dependent quantities that *do not* constitute existence claims.

19.2.1 QFT/GR

Quantum field theory (QFT).

- *Resonance/particle.* A particle-like excitation corresponds to a pole of a gauge-invariant two-point function or of the S-matrix at $p^2 = m_\star^2 - i m_\star \Gamma_\star$. The pair $(m_\star, \Gamma_\star)$ is invariant; the vector potential A_μ or a field component in a particular gauge is not. **Bridge:** a Breit–Wigner peak in a calibrated cross section or spectral density, with stability under binning and background subtraction. **Gates:** positivity of the Källén–Lehmann spectral density (S), causal support of response functions (G), complete positivity of reduced dynamics for open probes (TQ), and non-negative effective linewidths (D).

- *Charge and flux.* Electric charge is defined by a surface integral of the gauge-invariant field strength (Gauss law). **Bridge:** calibrated flux through a sensor array; **non-**

example: the value of a scalar potential in a chosen gauge.

- *Confinement signal.* The area law of a large Wilson loop, $\langle W(C)\rangle \sim \exp[-\sigma A(C)]$, is a gauge-invariant statement about the vacuum. **Bridge:** string tension inferred from heavy–quark spectroscopy and lattice determinations; report stability with respect to smearing windows and finite-volume effects.

General relativity (GR).

- *Curvature and horizons.* Coordinate singularities are not existence claims. Invariants such as the Kretschmann scalar $K = R_{\mu\nu\rho\sigma}R^{\mu\nu\rho\sigma}$ or the signs of null expansions on a trapped surface carry content. **Bridge:** gravitational-wave ringdown frequencies (quasi-normal mode spectra), lensing invariants, or redshift measurements tied to worldlines.

- *Gravitational waves.* What exists physically is the *geodesic deviation* encoded in components of the Riemann tensor along the detector worldlines, not a metric component in a chosen gauge. **Bridge:** calibrated differential arm length (interferometer strain) with a stated response function.

- *Mass/energy notions.* ADM or Bondi mass are coordinate-invariant but *quasi-local*/asymptotic descriptors; their use requires an explicit DoV statement (asymptotic flatness regime, extraction radius, numerical gauge conditions).

Table 19.2: Examples: from claim to gauge-invariant descriptor and pointer proxy.

Claim	**Invariant descriptor** D	**Pointer proxy / bridge**
"There is a new resonance"	S-matrix pole $(m_\star, \Gamma_\star)$	Peak in calibrated $\sigma(E)$ or spectral density after windowing; finite audit on binning.
"Charge enclosed"	$Q = \frac{1}{4\pi}\int_{S^2} \vec{E}\cdot d\vec{S}$	Calibrated flux through sensor array; geometry reported.
"Black hole present"	Trapped surface (null expansions $<$ 0); ringdown spectrum	GW strain modes (quasi-normal); lensing invariants; redshift timing.
"GW detected"	Riemann-induced geodesic deviation	Differential arm length (interferometer) with response function and noise budget.

Rule 19.2: Non-examples (do not reify)

- A component of A_μ or of $g_{\mu\nu}$ in a convenient gauge.
- "Superluminal recession speed" in cosmological coordinates (coordinate artefact).
- A divergence that disappears under a coordinate change or fails stability-at-finite.

Publishing posture. When asserting existence, authors should attach a one-paragraph *gate line* (S/G/TQ/D), a DoV declaration with units, and a pointer bridge. Competing descriptions are welcome, but the invariant D and the bridge must be *common*, so that disagreement is about physics, not about gauges.

19.3 Probability as Dissipation

Idea. In open dynamics the apparatus is not a spectator: it provides *channels* through which quanta and information flow. Let the reduced system evolve under a GKSL generator with jump operators collected into disjoint *readout blocks* (channels) $\{L_{r,k}\}_k$ labelled by $r = 1, \ldots, R$. Each block is wired to a distinct, calibrated counter (photodiode pixel, calorimetric pad, spin-flip register, *etc.*). Long-time counting then yields stable *shares* of dissipation into those channels. These shares are what we call "probabilities".

Definition 19.1: Dissipative share.

For a GKSL master equation

$$\dot{\rho} \;=\; -\,i[H,\rho] \;+\; \sum_{r=1}^{R}\sum_{k}\Big(L_{r,k}\rho L_{r,k}^{\dagger} - \tfrac{1}{2}\{L_{r,k}^{\dagger}L_{r,k},\rho\}\Big), \tag{19.1}$$

define for each channel the positive operator $E_r := \sum_k L_{r,k}^{\dagger}L_{r,k}$ and the steady state ρ_* of the dynamics (when it exists and is mixing on the reporting timescale). The *dissipative share* of channel r is

$$p_r \;=\; \frac{\mathrm{Tr}\Big(E_r\,\rho_*\Big)}{\sum_{s=1}^{R}\mathrm{Tr}\Big(E_s\,\rho_*\Big)}\,. \tag{19.2}$$

Operationally, if $N_r(T)$ counts the registered quanta in channel r during time window $[0,T]$, then under mild ergodicity conditions $p_r = \lim_{T\to\infty} N_r(T)/\sum_s N_s(T)$, with finite–$T$ corrections controlled by mixing and dead-time audits.

Why this earns "probability". (i) $E_r \geq 0$ guarantees spectral safety for the reported counts; (ii) the channel decomposition respects causality within g_{eff} (no superluminal

signalling between counters); (iii) complete positivity of (19.1) ensures non-negativity of all marginals for subsystems; (iv) the net entropy production rate $\sigma \geq 0$ provides the arrow of time (non-negative dissipation). Thus S/G/TQ/D are baked into (19.2).

Criterion 19.2: Admissibility for "probability as dissipation".

Use (19.2) when: (a) the counter partition $\{L_{r,k}\}$ is fixed and calibrated; (b) a mixing or stationary regime on the reporting window exists ($T \gg \tau_{\text{mix}}$); (c) stability at finite is documented with respect to binning, thresholds, and background subtraction; (d) uncertainty budgets (statistical and systematic) are reported.

title=How to report

State $(H, \{L_{r,k}\})$, the DoV (resolution, window, thresholds), the bridge from E_r to device counts, and the S/G/TQ/D gate line. Provide a stability plot showing $p_r(T)$ settling within the budget on the chosen window.

19.3.1 Open Models

We illustrate (19.2) with minimal, publishable templates.

Model A: Two-level system with two radiative exits. Let $H = \frac{\omega_0}{2}\sigma_z$ and take radiative jumps into two distinct, spatially separated detectors:

$$L_1 = \sqrt{\kappa_1}\,\sigma_-, \qquad L_2 = \sqrt{\kappa_2}\,\sigma_-.$$

Thermal re-excitation at rate λ closes the cycle (e.g. weak pump or hot background), $L_\uparrow = \sqrt{\lambda}\,\sigma_+$. The steady excited population is

$$\pi_e \;=\; \frac{\lambda}{\lambda + \kappa_1 + \kappa_2}$$

and the (long-time) photon fluxes into the two detectors are $J_r = \kappa_r\,\pi_e$. Hence the dissipative shares are

$$p_1 \;=\; \frac{J_1}{J_1 + J_2} \;=\; \frac{\kappa_1}{\kappa_1 + \kappa_2}, \qquad p_2 \;=\; \frac{\kappa_2}{\kappa_1 + \kappa_2}. \tag{19.3}$$

Equation (19.3) equals $\mathrm{Tr}(E_r\rho_*)$ with $E_r = \kappa_r\,\sigma_+\sigma_-$. *Bridge:* counts per window in each pixel; *stability-at-finite:* verify independence from bin width once $T \gg \tau_{\text{mix}}$; *gate line:* $E_r \geq 0$ (S), radiative light cone (G), GKSL (TQ), non-negative linewidths (D).

Model B: Damped oscillator and calorimetric bins. For a mode $H = \omega a^\dagger a$ coupled to a broadband bath:

$$\dot{\rho} = -i[H, \rho] + \gamma_\downarrow\Big(a\rho a^\dagger - \tfrac{1}{2}\{a^\dagger a, \rho\}\Big) + \gamma_\uparrow\Big(a^\dagger \rho a - \tfrac{1}{2}\{a a^\dagger, \rho\}\Big).$$

Let the calorimeter register energy packets $\hbar\omega$ leaving to two absorbers with couplings $\gamma_\downarrow^{(1)}$ and $\gamma_\downarrow^{(2)}$ (so $L_1 = \sqrt{\gamma_\downarrow^{(1)}}\, a$, $L_2 = \sqrt{\gamma_\downarrow^{(2)}}\, a$). In the stationary regime the mean outgoing quanta are proportional to those couplings; the shares are again $p_r = \mathrm{Tr}(E_r \rho_*)/\sum_s \mathrm{Tr}(E_s \rho_*)$ with $E_r = \gamma_\downarrow^{(r)} a^\dagger a$. The FDT appears as the relation between $\gamma_\uparrow/\gamma_\downarrow$ and the bath temperature; entropy production is $\sigma = (\gamma_\downarrow - \gamma_\uparrow)\, \hbar\omega\, \langle a^\dagger a\rangle / T_{\mathrm{bath}} \geq 0$ on the reporting window.

Model C: Classical limit / birth–death counting. Consider a classical Markov chain with states i and edges grouped into exclusive readout classes r. The master operator W has positive off-diagonal rates $w_{j\leftarrow i}^{(r)}$. The stationary distribution π obeys $W^\top \pi = 0$, and the long-time fraction of events in class r is

$$p_r = \frac{\sum_i \pi_i \sum_j w_{j\leftarrow i}^{(r)}}{\sum_s \sum_i \pi_i \sum_j w_{j\leftarrow i}^{(s)}}.$$

This is the classical counterpart of (19.2) with E_r the edge-incidence functional. The mapping clarifies that "probability" is a *rate share* in steady flow.

Rule 19.3: Open–model template (author checklist)

1. Specify $(H, \{L_{r,k}\})$ and the counter partition; state device calibration.
2. Prove/argue mixing on the reporting window; give ρ_* or the effective ergodic average used.
3. Report $p_r(T)$ vs T to document stability at finite; quote the chosen T and the residual drift.
4. Provide an uncertainty budget (shot noise, background, dead time).
5. Gate line: spectral safety of E_r; causal cone for the channels; CP of the reduced map; non-negative entropy production on the stated window.

Remark (beyond steady state). When only preasymptotic windows are accessible, one may use *tilted* (χ-deformed) generators to characterise finite-time counting statistics. The zero–tilt slope gives mean shares, and large–deviation bounds provide honest finite–T error bars; all reporting must keep the DoV explicit.

19.4 Internal Time

Motivation. There is no external master clock in physics. What we actually read are *devices* that tick; their states provide monotone parameters against which other processes are compared. The content of a time claim is therefore *relational*: differences between clock readings, frequency ratios, and phase lags along specified, causally admissible worldlines. This section formalizes that stance and gives a reporting template.

Principle 19.2: No absolute time.

Only differences and rates defined by physical clocks are meaningful. Reparametrizations of an external parameter $t \mapsto f(t)$ must not change any reported observable; device–level reports use *differences* Δ and *ratios* that are invariant under such changes.

Definition 19.2: Internal clock.

A system C with a calibrated pointer map R_C whose reading $C(\tau)$ changes monotonically on its domain of validity (DoV). The *time differences* it defines on the window $[\tau_1, \tau_2]$ are $\Delta C := C(\tau_2) - C(\tau_1)$; its *rate* against another clock D is the dimensionless ratio dC/dD (understood as a finite difference $\Delta C/\Delta D$ on the reporting window).

Definition 19.3: Internal-time claim.

A statement about dynamics is device-facing if it is expressed in terms of (i) differences Δ of calibrated readings, or (ii) ratios of frequencies or phases measured by two clocks, and if it passes the S/G/TQ/D gates within the effective cone g_{eff} relevant to the measurement.

Criterion 19.3: Admissibility for internal-time reports.

Declare: (a) the clocks used (physics and calibration); (b) the synchronization protocol and light-cone assumptions; (c) the DoV and window length; (d) the invariance of the reported quantity under $t \mapsto f(t)$ and unit changes; (e) uncertainty budgets (phase noise, dead time, transport delays).

Reparametrization invariance (worked check). Let $C(\tau)$ and $D(\tau)$ be two monotone clock readings on the same record. Then the observable rate is

$$\frac{dC}{dD} = \frac{dC/d\tau}{dD/d\tau},$$

which is unchanged under any smooth monotone change $\tau \mapsto f(\tau)$. Hence frequency *ratios* and phase *differences* are intrinsically gauge-free notions of time.

Examples. *Ramsey spectroscopy:* after free evolution for a window T_D defined by a reference clock D, the phase of system S advances by $\Delta\phi_S = (\omega_S - \omega_{\text{ref}})\, T_D$. The reportable quantity is either $\Delta\phi_S$ itself or the frequency ratio $r = \omega_S/\omega_D$.

Clock ratios: optical clocks A, B yield $r = \nu_A/\nu_B$ which is unitless and stable under device rescaling. Transport and gravitational redshift corrections are handled in the synchronization protocol.

title=GR redshift as an internal-time relation

For stationary fields with $g_{00} \approx -(1 + 2\Phi/c^2)$, two colocated identical clocks at potentials Φ_A, Φ_B satisfy on small baselines

$$\frac{\nu_A}{\nu_B} \approx 1 + \frac{\Phi_A - \Phi_B}{c^2}.$$

This is a *ratio* (difference observable) plus a causal statement about how the two clocks were compared (exchange along future-directed paths inside g_{eff}).

Open dynamics view. When clocks are noisy, we model (S, C) jointly with a GKSL generator and read phases/energies at the pointers. Reported rates dC/dD must be stable at finite window sizes and insensitive to admissible coarse-graining.

Rule 19.4: Internal-time report (author checklist)

1. Identify the clocks (physics, pointer maps, calibration).
2. Specify the synchronization protocol and causal paths.
3. State the reported quantity (difference or ratio) and show invariance under $t \mapsto f(t)$ and unit changes.
4. Give the DoV and window; document stability-at-finite and the uncertainty budget.
5. Gate line: spectral safety of readouts; causality within g_{eff}; CP of reduced maps for noisy clocks; non-negative dissipation on the window (arrow of time).

19.4.1 Difference Observables

Idea. In practice we never quote an absolute phase or time; we quote *differences* between two worldlines and two devices. We formalize a general construction that covers phase, frequency, and proper-time comparisons.

Definition 19.4: Difference observable.

Given two calibrated readouts R_A, R_B producing series $A(\tau_A)$ and $B(\tau_B)$ along specified worldlines, and a synchronization map that assigns admissible pairs of events (τ_A, τ_B), a difference observable is a functional

$$\Delta_{A\leftrightarrow B} := \mathcal{F}\Big(A(\tau_A), B(\tau_B)\Big),$$

typically $\Delta A - \Delta B$, a phase lag $\Delta\phi_A - \Delta\phi_B$, or a ratio dA/dB. The definition is required to be invariant under monotone reparametrizations of the individual records and under unit rescalings.

Construction patterns.

1. *Beat-note construction.* Mix two periodic signals and read the envelope frequency: $r = \nu_A/\nu_B$ from the beat $\nu_{\text{beat}} = |\nu_A - \nu_B|$. Robust and unitfree.

2. *Two-way-time transfer.* Use exchanged pulses to define a symmetric pairing of events and report $\Delta C - \Delta D$ with calibrated path delays; causality within g_{eff} is explicit.

3. *Proper-time differences.* For worldlines γ_A, γ_B in a stationary background, report $\Delta\tau_A - \Delta\tau_B = \int_{\gamma_A} \sqrt{-g_{\mu\nu}dx^\mu dx^\nu} - \int_{\gamma_B} \sqrt{-g_{\mu\nu}dx^\mu dx^\nu}$, approximated by frequency ratios of physical clocks transported along those paths.

title=Typical failure modes to avoid

Unit-dependent "absolute times"; mixing of coordinate time with device time; comparing non-synchronized events; hidden dependence on t-binning; neglecting dead-time or transport delays; using model phases without a pointer bridge.

Reporting example (Ramsey pair). Two qubits A, B interrogated for T with the same reference yield phases $\Delta\phi_A = \delta_A T$, $\Delta\phi_B = \delta_B T$. The difference observable is $\Delta\phi_A - \Delta\phi_B = (\delta_A - \delta_B)\,T$ and the ratio $r = \delta_A/\delta_B$. Provide the synchronization description (common pulses), the window T, phase-noise budget, and a stability plot vs T.

Takeaway. "Internal time" means that clocks are part of the model and of the audit trail. The only quantities that survive gauge changes and unit conventions are differences and ratios tied to calibrated devices and causal comparisons. That is the level at which time enters this volume's measurement pipelines.

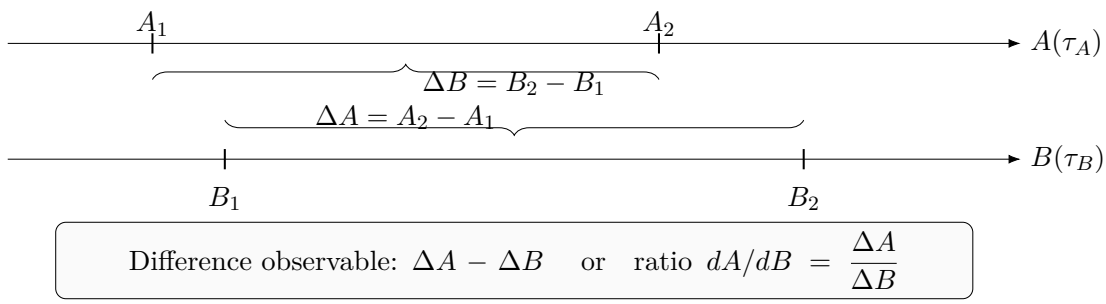

Figure 19.1: Internal time: report differences and ratios between calibrated clocks along admissible worldlines.

Chapter 20

Notation and Conventions

This appendix fixes typography, units, index rules, transform conventions, and reporting style used across the volume. Its purpose is practical: to make every symbol *auditable* (units, DoV, gates) and every equation *portable* across gauges and coarse–graining choices.

20.1 Typography and naming

Rule 20.1: Objects and their glyphs (house style)

Scalars (physical)	Roman x, T, Γ, $m_\star$, ρ_{ii} with units (SI by default).
Vectors/fields	Bold roman $\mathbf{E}$, $\mathbf{B}$, $\mathbf{v}$; components carry indices.
Matrices/operators	Upright roman H, $L_{r,k}$, E_r, U, ρ; spectra via $\mathrm{Spec}(\cdot)$.
States (QM)	Bras/kets $\langle\phi\vert$, $\vert\psi\rangle$; density operator ρ.
Random/counting vars	Sans serif when helpful $\mathsf{N}_r(T)$; otherwise roman with context.
Transforms	$\mathcal{F}\{x\}(\omega)$ (Fourier), $\mathcal{L}\{x\}(s)$ (Laplace).
Invariants	Named explicitly (e.g. "Kretschmann scalar K", "Wilson loop $W(C)$").

Reserved macros. We use Tr (trace), Spec (spectrum), GKSL (generator class), CP (complete positivity), and the gate tags S, G, TQ, D. The device–inferred cone/metric is g_{eff}.

20.2 Indices, tensors, and sign conventions

- Greek indices $\mu, \nu, \ldots$ for spacetime components; Latin i, j, k for spatial, modes, or channels (context clarified at first use).
- Metric signature $(-,+,+,+)$ unless stated otherwise. Raising/lowering by $g_{\mu\nu}$ (or by g_{eff} when explicitly working with effective cones).
- Riemann tensor convention: $R^{\rho}{}_{\sigma\mu\nu} = \partial_\mu \Gamma^{\rho}_{\nu\sigma} - \partial_\nu \Gamma^{\rho}_{\mu\sigma} + \cdots$. Kretschmann scalar $K = R_{\mu\nu\rho\sigma} R^{\mu\nu\rho\sigma}$.
- Levi–Civita $\epsilon^{0123} = +1$; dual field $\tilde{F}^{\mu\nu} = \frac{1}{2}\epsilon^{\mu\nu\rho\sigma} F_{\rho\sigma}$.

title=Coordinate vs invariant

Any chart/foliation/lapse–shift choice is descriptive. Claims are made in invariants (curvature scalars, holonomies, spectra, cone order) or in entire equivalence classes that reach a pointer; see Chapter 18.

20.3 Units, dimensions, constants

Default. SI units via `siunitx`. Every reported quantity carries units or is a dimensionless *ratio* tied to calibrated units.

Natural units. When useful we set $\hbar = k_B = 1$; c is kept explicit in GR contexts unless noted. Any departure from SI is flagged inline (box or parenthetical) and never assumed silently.

Table 20.1: Frequent symbols with units and quick meaning.

Symbol	Units	Meaning	Where used
H	J or $\hbar\,\mathrm{s}^{-1}$	Hamiltonian	Dynamics
$L_{r,k}$	$\mathrm{s}^{-1/2}$	Jump operators (channel r)	Open systems
$E_r = \sum_k L^\dagger_{r,k} L_{r,k}$	s^{-1}	Channel intensity operator	Shares
$S_{YY}(\omega)$	$(\text{unit})^2/\text{Hz}$	One–sided PSD	Spectral reports
$m_\star, \Gamma_\star$	eV, eV	Pole mass and width	QFT existence
K	m^{-4}	Curvature invariant	GR audits
g_{eff}	—	Effective causal cone/metric	Causality checks

Dimensional checks. We encourage inline checks: $[S_{YY}] = [Y]^2/\text{Hz}$; $[E_r] = \mathrm{s}^{-1}$; $[\Gamma_\star] = \mathrm{s}^{-1}$ or $\text{eV}/\hbar$; report conversions when needed.

20.4 Fourier/Laplace and spectral conventions

Fourier transform. For a (real) stationary process $Y(t)$,

$$\mathcal{F}\{Y\}(\omega) = \int_{-\infty}^{\infty} Y(t)\, e^{-i\omega t}\, dt, \qquad S_{YY}(\omega) = \mathcal{F}\{C_{YY}\}(\omega) \geq 0 \text{ (one–sided on the DoV).}$$

We use one–sided PSDs in figures; window/taper/estimator is always stated.

Laplace/Green's functions. Retarded response obeys causal support: $G_R(t < 0) = 0$. Front velocity (isochrons) implements G; group–velocity exceedance does not constitute signalling.

Criterion 20.1: PSD admissibility

Quote PSDs only with estimator, window T, binning, and a statement that the Toeplitz estimate is positive semidefinite (spectral safety, S).

20.5 Open dynamics and GKSL sign conventions

We use the standard GKSL form

$$\dot{\rho} = -\, i[H, \rho] + \sum_{r,k} \Big(L_{r,k} \rho L_{r,k}^{\dagger} - \tfrac{1}{2}\{L_{r,k}^{\dagger} L_{r,k}, \rho\} \Big),$$

with readout *channels* labelled by r. The operators $E_r = \sum_k L_{r,k}^{\dagger} L_{r,k} \geq 0$ define calibrated intensities tied to counters. Born appears as *stable dissipative shares* p_r; see Chapter 9 and section 19.3.

20.6 Probability and statistics (reporting)

Rule 20.2: Uncertainty and stability (house rules)

1. Quote statistical (counting) and systematic (calibration/background) errors separately when possible; otherwise provide a combined budget and method.
2. When a limit is used ($T \to \infty$, continuum, plane waves), report a *stability-at-finite* plot or table and a tolerance ε tied to device resolution.
3. Convergence language: "stabilizes within ε on window T" beats bare asymptotic notation.

Symbols. We use $E[\cdot]$ for expectation, $\mathrm{Var}[\cdot]$ variance, $\mathrm{Cov}[\cdot,\cdot]$ covariance; "share vectors" $\mathbf{p}$ are probabilities (non-negative, sum to 1) inferred from counts.

20.7 Domains of validity (DoV), gates, and badges

title=Badge lines (equations/figures)

After important equations and in figure captions we append a *badge line*: **Units**; **DoV** (band/scale/window and end-of-validity triggers); **Gauge** status (invariant or choice flagged); **Finite audit** (knobs and tolerance); **Gates** touched (S/G/TQ/D); **Bridge** (model→pointer map and calibration ref).

Criterion 20.2: Admissible claim

A sentence advertising a physical effect is admissible iff it includes or directly references: Units, DoV, an invariant form, a finite-audit when limits are used, a gate line, and a pointer bridge. Missing items ⇒ the sentence is mathematical/descriptive.

20.8 Spectrum-first descriptors (object language policy)

"Object/particle" language is licensed only by a *spectral anchor*: shares $\{p_m\}$, peak sets/Qs, poles/residues with units, stability-at-finite, gate line, and a pointer bridge (Chapter 4). Without this, rewrite the claim in spectral terms.

20.9 Figures, tables, and captions

Figures. Axes labelled with units; error bars or shaded uncertainty; estimator, window, and binning declared for spectral plots; `\resizebox` used to avoid margin overflow.

Tables. Use `booktabs`; units in headings; significant figures consistent with uncertainty. Prefer "long-ish" captions that include the badge items.

20.10 Citations and reproducibility (practical)

Citations. Where bibliographic tools are enabled, we prefer author–year and compact citestyle; when not, include inline short references. Equations and data should be citable by label and figure number, not only by page.

Table 20.2: Abbreviations used in the volume.

Abbrev.	Expansion	Context
PSD	Power spectral density	Spectral reports
DoV	Domain of validity	Acceptance metrics
CP/CPTP	(Completely) positive, trace-preserving	Open maps
QND	Quantum non-demolition	Measurement models
RG	Renormalization group	Scale transformations
EFT	Effective field theory	Coarse–grained models
AB	Aharonov–Bohm	Gauge invariants
ADM/Bondi	Asymptotic mass notions	GR reports
EP	Entropy production	Dissipative arrows

Reproducibility. Whenever feasible, attach data/estimator scripts (windowing, binning), calibration files, and random seeds. If only processed data can be released, describe the pipeline sufficient for independent reproduction at the DoV.

20.11 Worked micro-templates

Rule 20.3: Equation badge (fill and paste)

Units: $[\cdot] = \cdots$. **DoV:** band/scale/window; end-of-validity. **Gauge:** invariant / descriptive choice. **Finite audit:** knobs + tolerance $\varepsilon = \cdots$. **Gates:** S (...), G (...), TQ (...), D (...). **Bridge:** model→pointer (calibration ref).

Rule 20.4: Figure caption badge (fill and paste)

Units, **DoV**, **Invariant**, **Finite audit**, **Gates**, **Bridge**.

These conventions exist to keep mathematics powerful *without* letting it slip into ontology. They align with the volume's central themes: gauge realism, stability at finite resolution, spectrum-first reporting, and the S/G/TQ/D gate discipline.

Chapter 21

Mini-Derivations and Worked Examples

This appendix collects compact derivations and end–to–end worked examples that instantiate the book's workflow: invariants → stability-at-finite → gate line → pointer bridge. Each item is self-contained and ready to cite.

21.1 Mini-derivation 1 — PSD non-negativity at finite resolution

Proposition 21.1 (Bochner at finite)**.** *Let $Y(t)$ be a (wide-sense) stationary process with autocovariance $C(\tau) = E[Y(t)Y(t+\tau)]$. On a finite window with samples $Y_0, \ldots, Y_{N-1}$, the Toeplitz covariance matrix $T_{ij} = C(i-j)$ is positive semidefinite. Any periodogram formed as $\widehat{S}(\omega) = \mathbf{w}^* T \mathbf{w}$ with a taper $\mathbf{w}$ is non-negative. In the continuous limit, Bochner's theorem gives $S_{YY}(\omega) \geq 0$.*

Sketch. For any vector $\mathbf{z}$, $\mathbf{z}^* T \mathbf{z} = E\left|\sum_j z_j Y_j\right|^2 \geq 0$, hence $T \succeq 0$. Choosing $z_j = w_j e^{-i\omega j \Delta t}$ yields the tapered periodogram $\widehat{S}(\omega) = \mathbf{w}^* T \mathbf{w} \geq 0$. □

Badge: Units: $[S_{YY}] = [Y]^2/\text{Hz}$. DoV: window $T = N\Delta t$, taper/multitaper stated. Gates: S (PSD ≥ 0); G n/a; TQ n/a; D n/a. Bridge: analyzer calibration.

21.2 Mini-derivation 2 — Born as long-time dissipative shares

Consider one acquisition cycle $\mathcal{T} = (\sum_r \mathcal{I}_r) \circ \mathcal{E}_\Delta$, with instrument maps $\mathcal{I}_r(\rho) = \sum_k M_{r,k} \rho M_{r,k}^\dagger$ and $E_r = \sum_k M_{r,k}^\dagger M_{r,k}$. Between shots, $\mathcal{E}_\Delta$ is CP (open dynamics on the DoV).

Proposition 21.2 (Pipeline LLN)**.** *If $\mathcal{T}$ is mixing with unique steady state ρ_*, the empirical shares $s_r(N) = \frac{1}{N}\sum_{n=1}^{N} X_n^{(r)}$ (where $X_n^{(r)}$ flags outcome r at cycle n) satisfy*

$$s_r(N) \xrightarrow[N\to\infty]{a.s.} \mathrm{Tr}(E_r\rho_*), \qquad \sqrt{N}\Big(s_r(N) - \mathrm{Tr}(E_r\rho_*)\Big) \Rightarrow \mathcal{N}(0, \sigma_r^2).$$

Sketch. Conditioned on ρ_n, $E[X_n^{(r)}|\rho_n] = \mathrm{Tr}(E_r\rho_n)$. The martingale difference $X_n^{(r)} - \mathrm{Tr}(E_r\rho_n)$ obeys a LLN. Mixing implies $\rho_n \to \rho_*$, yielding the limit. A martingale CLT gives the Gaussian fluctuations with variance set by instrument and gap. □

Badge: Units: counts/shot. DoV: $N \gg 1$, mixing gap reported. Gates: S ($E_r \geq 0$), G (no-signalling channels), TQ (CP), D (non-negative linewidth/EP). Bridge: counter calibration.

21.3 Mini-derivation 3 — Stability-at-finite acceptance test

Let $Q = \lim_{\Lambda\to\infty} Q(\Lambda)$ be used in a claim.

Criterion 21.1: Finite Stability

License Q on the DoV if there exist Λ_0, ε such that for all $\Lambda \geq \Lambda_0$ and fractional steps $|\delta| \leq \Lambda/2$,

$$\frac{|Q(\Lambda+\delta) - Q(\Lambda)|}{\max\{1, |Q(\Lambda)|\}} < \varepsilon,$$

with ε tied to device tolerance and quoted. Otherwise, reformulate the claim or change the observable/model.

Badge: Units: as Q. DoV: $\Lambda \in [\Lambda_0, 2\Lambda_0]$. Gates: depends on context. Bridge: how Q maps to pointer.

21.4 Worked example A — UV observable with a hard cutoff

Setup. A scalar loop gives $O(\Lambda) = \int^{\Lambda} \frac{d^4k}{(2\pi)^4} \frac{f(k)}{(k^2+m^2)^2}$ for a device-facing observable O (e.g. slope of a calibrated cross-section). We test stability-at-finite.

Computation (schematic). For large Λ, $O(\Lambda) = a_0 + a_1 \log\frac{\Lambda}{\mu} + O(\Lambda^{-2})$. Define the renormalized $O_R(\mu) = O(\Lambda) - a_1 \log(\Lambda/\mu)$. Then $|O_R(\mu; \Lambda+\delta) - O_R(\mu; \Lambda)| = O(\Lambda^{-2})$ passes Criterion **??** on a declared window.

Badge. Units: as O. DoV: $\Lambda \in [1, 2]$ GeV (example), tolerance $\varepsilon = 0.02$. Gates: S (positivity of reported spectrum if applicable). Bridge: how O_R is read from the fit. Falsifier: drift $> \varepsilon$ under Λ shift.

21.5 Worked example B — Front vs group velocity (no–signalling)

Setup. Medium with retarded Green function $G_R(x, t)$ and dispersion $\omega(k)$. Group velocity $v_g = d\omega/dk$ can exceed $c_\star$; front velocity $v_f = \lim_{k\to\infty} \omega/k$ controls causal support.

Derivation (sketch). Causality implies $G_R(x, t) = 0$ for $|x| > v_f t$. A pump at $x = 0$ yields probe response at $x = L$ only for $t \geq L/v_f$. Thus any pointer change before L/v_f is ruled out (within noise), regardless of v_g.

Badge. Units: $[v_f]$ = m/s. DoV: band $[\omega_1, \omega_2]$; isochron detection. Gates: G (no-signalling), S (PSD ≥ 0). Bridge: pump–probe isochrons.

21.6 Worked example C — Resonance pole extraction (QFT)

Setup. Differential cross-section $\sigma(E)$ exhibits a bump near $E \simeq m_\star$.

Algorithm. (1) Subtract smooth background (calibrated polynomial/K–matrix). (2) Fit a Breit–Wigner form $\sigma(E) \propto \frac{\Gamma_\star^2/4}{(E-m_\star)^2+\Gamma_\star^2/4}$. (3) Check stability under bin-size change and fit window shifts. (4) Report $(m_\star, \Gamma_\star)$ with uncertainties.

Invariant. The S–matrix pole $E_\star = m_\star - i\Gamma_\star/2$ is gauge/field–redefinition invariant; components of a vector potential are not.

Badge. Units: eV. DoV: window $\pm 5\Gamma_\star$; binning audit. Gates: S (positive spectral density), TQ (CP if reduced). Bridge: detector response.

21.7 Worked example D — GR ringdown as invariant spectroscopy

Setup. After a compact merger the strain is a sum of quasi-normal modes (QNMs): $h(t) = \sum_n A_n e^{-t/\tau_n} \cos(2\pi f_n t + \phi_n)$.

Invariant content. $\{(f_n, \tau_n)\}$ are the *poles* of the retarded Green function of the curvature perturbation; hence invariant under coordinate change.

Bridge. Fit peaks in the PSD and exponential decays in the time series with a detector response model; quote (f_n, τ_n) and their covariance.

Badge. Units: Hz, s. DoV: post–merger window; SNR threshold; calibration lines removed. Gates: G (causal propagation), S (PSD), D (positive damping). Falsifier: coordinate–dependent "feature" disappearing under invariant restatement.

21.8 Worked example E — Internal time via frequency ratios

Setup. Two clocks A, B produce readouts with carrier frequencies ν_A, ν_B. We measure the ratio $r = \nu_A/\nu_B$.

Beat-note derivation. Mix signals to obtain a beat $\nu_{\text{beat}} = |\nu_A - \nu_B|$. Then $r = \nu_A/\nu_B = 1 \pm \nu_{\text{beat}}/\nu_B$ (use the sign from prior). Under any reparametrization $t \mapsto f(t)$, r remains unchanged.

Badge. Units: dimensionless. DoV: averaging time T; phase–noise model. Gates: G (causal transfer), TQ (CP if quantum clocks), D (non-negative EP). Bridge: counter and mixer calibration.

21.9 Worked example F — Spectrum-first soliton audit (NLSE proxy)

Setup. A field obeys an effective NLSE with a localized packet observed on a window.

Spectral anchor. Compute the share vector p_m in a calibrated mode basis (or the peak set in $S_{YY}(\omega)$). Localization via inverse participation ratio

$$\mathrm{IPR} = \sum_m p_m^2 \quad (\text{large IPR} \Rightarrow \text{localized}).$$

Audit stability under (i) window growth, (ii) injected loss/gain within the DoV, (iii) small dispersion changes.

Badge. Units: p_m dimensionless; frequencies in Hz. DoV: band and acquisition time. Gates: S (PSD ≥ 0), D (linewidths ≥ 0), G (fronts within g_{eff}). Bridge: spectrum analyzer/camera calibration. Falsifier: IPR collapse under admissible coarse-graining.

21.10 Worked example G — CP of GKSL maps (Spohn inequality sketch)

Setup. $\dot{\rho} = \mathcal{L}(\rho)$ in GKSL form. Let ρ_{ss} be a full-rank stationary state ($\mathcal{L}^\dagger(\rho_{\mathrm{ss}}) = 0$).

Claim. Entropy production $\sigma(t) = -\frac{d}{dt} D(\rho_t \| \rho_{\mathrm{ss}}) \geq 0$, where $D(\cdot\|\cdot)$ is quantum relative entropy. This implies D and is consistent with TQ.

Sketch. Spohn's inequality for GKSL generators gives $\frac{d}{dt}\,\mathrm{Tr}[\rho_t(\log \rho_t - \log \rho_{\mathrm{ss}})] \leq 0$. Hence $D(\rho_t \| \rho_{\mathrm{ss}})$ decreases, defining $\sigma \geq 0$.

Badge. Units: nats/s. DoV: validity of time-local GKSL. Gates: TQ, D. Bridge: estimate σ via heat flows and effective temperatures.

21.11 Worked example H — Coordinate artefact audit (GR horizon)

Setup. Metric component g_{tt} diverges at $r = r_h$ in a chart.

Invariant diagnosis. Compute curvature invariants (e.g. K). If finite at r_h, the divergence is coordinate. Switch to Eddington–Finkelstein/Kruskal coordinates; physical pointers (time delay, redshift) remain finite across r_h.

Badge. Units: K in m^{-4}. DoV: near-horizon regime. Gates: G (causal order), S (PSD for responses). Bridge: lensing/time-delay calibration. Falsifier: invariant blow-up (then it's end-of-validity; replace the model).

How to use this appendix. Each mini–derivation/example can be copied with its badge line into the main text or supplementary notes. The structure ensures that claims are invariant, stable at finite resolution, pass the gates, and reach a calibrated pointer—precisely the discipline advocated in the volume.

Chapter 22

Gate Checklists and Templates

This appendix is the practical handbook for applying the four consistency gates—S, G, TQ, D—across the volume. It provides (i) a one–page overview, (ii) operational tests per gate, (iii) ready–to–paste gate cards and report templates, and (iv) a few filled examples that demonstrate the level of evidence expected. Use it when drafting claims, reviewing sections, or designing figures.

22.1 One–page overview (what each gate asks for)

title=How to read a gate line

A *gate line* is a single sentence per gate, tied to the stated DoV, that references the concrete check performed (matrix eigenvalues; isochrons; Choi; entropy balance) and where its evidence lives (figure/table/appendix).

22.2 Operational tests by gate

S — Spectral Safety.

S1. *PSD positivity (finite window).* Form the tapered/multitaper estimate $\widehat{S}(\omega) = \mathbf{w}^* T \mathbf{w}$ where T is the Toeplitz covariance matrix. Check $\widehat{S}(\omega) \geq 0$ on the DoV (see Prop. 21.1).

S2. *Observable positivity.* For any reported covariance/Gram matrix C of observables, verify $\lambda_{\min}(C) \geq -\epsilon$ with ϵ tied to noise.

S3. *Field–theory spectral density.* When quoting spectral representations, require the Källén–Lehmann weight $\rho(\mu^2) \geq 0$ on the DoV; negative weight signals model/fit failure.

Table 22.1: The four gates: meaning, minimal evidence, common pitfalls.

Gate	Minimal evidence (device–facing)	Common pitfalls
S (Spectral Safety)	Non–negativity of spectra/-variances; PSD ≥ 0; positive semidefinite covariance/Choi where applicable; absence of ghost/negative–norm modes.	Using raw periodograms without PSD tapers; negative fitted variances; reporting "modes" with negative weight; unphysical sign in Källén–Lehmann density.
G (Causality within g_{eff})	Front/isochron maps consistent with the effective cone; microcausality or no–signalling tests; retarded support for Green functions.	Confusing group–velocity exceedance with signalling; coordinate artefacts; ignoring path delays and synchronization in comparisons.
TQ (Complete Positivity)	Choi matrix ≥ 0; trace preservation; GKSL form for time–local generators on the DoV; CP for reduced dynamics under coarse–graining.	Verifying positivity only on pure states; using non–CP effective maps after truncations; neglecting ancilla tests.
D (Non–Negative Dissipation)	Non–negative linewidths/absorption; entropy production rate $\sigma \geq 0$ on the reporting window; fluctuation–dissipation consistency.	Inferring negative damping from under–resolved fits; mixing reversible work with heat; declaring σ without specifying baths/temperatures.

G — Causality within g_{eff}.

G1. *Front test (pump–probe).* Map isochrons; require no pointer change before the front time $t_f = L/v_f$ within the uncertainty budget.

G2. *Microcausality/no–signalling.* In QFT, check vanishing commutators for spacelike separations on the DoV; in devices, verify intervention on channel A does not change B's statistics outside g_{eff}.

G3. *Retarded support.* Report that fitted/used Green functions G_R vanish for $t < 0$ (within tolerance) on the band used.

TQ — Complete Positivity.

TQ1. *Choi test (discrete map).* For a channel Φ, form the Choi matrix $J = (\Phi \otimes \mathbb{I})(|\Omega\rangle\langle\Omega|)$ with $|\Omega\rangle = \sum_i |ii\rangle$; require $J \succeq 0$ and $\mathrm{Tr}_{\text{out}}\, J = \mathbb{I}$ (trace preservation).

TQ2. *GKSL identification (time–local).* Exhibit H and $\{L_{r,k}\}$ such that $\dot{\rho} = -i[H,\rho] + \sum_{r,k}\left(L_{r,k}\rho L_{r,k}^\dagger - \frac{1}{2}\{L_{r,k}^\dagger L_{r,k}, \rho\}\right)$ reproduces the reduced dynamics on the DoV.

TQ3. *Coarse–graining safety.* If you trace out subsystems or bin in time, re–check CP of the induced map on the remaining degrees of freedom.

D — Non–Negative Dissipation (arrow of time).

D1. *Entropy production (quantum).* With stationary reference ρ_{ss}, report $\sigma(t) = -\frac{d}{dt} D(\rho_t \| \rho_{\mathrm{ss}}) \geq 0$ (Spohn inequality) on the window.

D2. *Linewidths/absorption.* Fit positive linewidths and non–negative absorption; negative values beyond uncertainty are inadmissible on the stated DoV.

D3. *FDT consistency.* Where applicable, check fluctuation–dissipation relations between noise spectra and response; inconsistencies indicate end–of–validity or mis–calibration.

22.3 Gate cards and report templates

Rule 22.1: Universal Gate Card (attach per claim/figure)

Claim (invariant, with units): .....
Domain of validity (DoV): band/scale/window; end–of–validity triggers .....
Gauge status: invariant statement; descriptive choices flagged .....
Limit audit: finite knobs (window/grid/cutoff) and tolerance $\varepsilon =$ .....
Gates: S (what was checked and where), G (front/no–signalling evidence), TQ (Choi/GKSL), D (EP/linewidths) .....
Bridge (model → pointer): calibration, estimator, figure/table ref .....
Falsifier: pointer–level signature that would refute the claim .....

Rule 22.2: Per–gate mini–cards (paste as needed)

S. PSD/covariance positivity; eigenvalues $\lambda_{\min} \geq -\epsilon$ (ref: Fig. # / Tbl. #).
G. Isochrons: no response before t_f; microcausality/no–signalling test passed (ref: Fig. #).
TQ. Choi $J \succeq 0$; $\mathrm{Tr}_{\mathrm{out}}\, J = \mathbb{I}$; GKSL form on DoV (ref: App. #).
D. $\sigma \geq 0$ on window; linewidths ≥ 0; FDT consistent (ref: Tbl. #).

title=Reviewer checklist (quick lint)

Units? DoV? Invariant form? Finite audit? S/G/TQ/D line? Bridge?
If any is missing, the sentence is *mathematical/descriptive*, not a physical claim.

22.4 Filled examples (what a passing card looks like)

Example 1 — S: interferometer PSD. *Claim.* The strain PSD $S_{hh}(\omega)$ around 150 Hz stays below $1 \times 10^{-23}\,\sqrt{\mathrm{Hz}}$ on the observing run.

DoV. Band 120 Hz to 180 Hz; window $T = 128\,\mathrm{s}$; 5–tap multitaper.
Checks. Toeplitz PSD $\succeq 0$; no negative bins beyond 2σ; calibration lines masked.
Gates. S pass (PSD), G n/a, TQ n/a, D n/a.
Bridge. Analyzer chain (ver. X.Y); calibration file CAL–2025–10.
Falsifier. Negative PSD bins beyond budget under the same estimator.

Example 2 — G: anomalous dispersion, no–signalling. *Claim.* Despite $v_g > c_\star$ on 1.8 GHz to 2.0 GHz, no signalling occurs.
DoV. Distance $L = 10\,\mathrm{m}$; isochron resolution 1 ns.
Checks. Pump–probe: no pointer change before $t_f = L/v_f$ within noise; front velocity from G_R matches g_{eff}.
Gates. G pass; S PSD ≥ 0; others n/a.
Bridge. Time–tagging scope + calibrated trigger offsets.
Falsifier. Reproducible pre–front response > noise.

Example 3 — TQ: qubit amplitude damping channel. *Map.* $\Phi(\rho) = K_0\rho K_0^\dagger + K_1\rho K_1^\dagger$ with $K_0 = \left(\begin{smallmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{smallmatrix}\right)$, $K_1 = \left(\begin{smallmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{smallmatrix}\right)$, $\gamma \in [0,1]$.
Choi. $J = (\Phi \otimes \mathbb{I})(|\Omega\rangle\langle\Omega|)$ has eigenvalues $\{1-\gamma,\ \gamma,\ 0,\ 0\} \geq 0$; $\mathrm{Tr}_{\mathrm{out}}\, J = \mathbb{I}$ (TP).
DoV. Gate time 20 ns; drift $< 1\,\%$.
Gates. TQ pass (CP/TP); D pass (linewidths ≥ 0).
Bridge. State/process tomography with calibrated SPAM errors.

Example 4 — D: entropy production for a driven TLS. *Model.* GKSL with upward/downward rates $\gamma_\uparrow, \gamma_\downarrow$ at bath T.
EP. $\sigma = (\gamma_\downarrow - \gamma_\uparrow)\frac{\hbar\omega_0}{T}\, p_e \geq 0$ on the window (measured heat current J_Q gives the same sign).
DoV. Weak driving; $T_{\mathrm{window}} \gg 1/\gamma_\Sigma$.
Gates. D pass; TQ pass (GKSL); S pass (rates ≥ 0); G n/a.
Bridge. Calorimetric power and excited–state fraction p_e via fluorescence.

22.5 Author templates by domain (paste and fill)

Rule 22.3: GR claim (template)

Invariant. Curvature/holonomy: .....
DoV. Region, coordinates (descriptive), extraction radius: .....
Finite audit. Resolution; gauge–insensitivity report: .....
Gates. G: fronts/causal order ; S: PSD of response ; TQ: n/a; D: damping signs
Bridge. Time delay, lensing, ringdown spectra (calibrated): .....

Rule 22.4: QFT claim (template)

Invariant. S–matrix pole / spectral residue: .....
DoV. Energy window and binning; scheme independence check: .....
Finite audit. Cutoff/fit–window stability: .....
Gates. S: Källén–Lehmann $\rho \geq 0$; G: retarded support ; TQ: CP if reduced ; D: widths ≥ 0 .....
Bridge. Cross section / spectral line calibration: .....

Rule 22.5: Open–quantum/measurement claim (template)

Invariant. POVM/Choi parameters, share vector $\mathbf{p}$: .....
DoV. Window T, dead time, thresholds: .....
Finite audit. $p_r(T)$ stabilization; bin sweep: .....
Gates. S: $E_r \geq 0$ ..; G: channel no–signalling ..; TQ: Choi/GKSL ..; D: $\sigma \geq 0$ ..
Bridge. Counter calibration; estimator details: .....

These checklists are intentionally terse and repetitive: they make it difficult to write claims that *sound* physical but are not tied to invariants, stability at finite resolution, the causal cone, complete positivity, and non–negative dissipation on a declared domain of validity.